Full Circle Magazine

Full Circle
THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY
ISSUE #97 - May 2015
KDE PLASMA 5
THE BEST KDE YET?
Full Circle Magazine is neither affiliated with, nor endorsed by, Canonical Ltd.

Issue #97

Get *unlimited access* to a cutting-edge technology and business library with **Apress Access!**

For $199

YOU GET:

- Unlimited access to Apress titles for a full year
- Instant access to each new Apress publication
- Compatibility with any device—desktop, laptop, or mobile
- Use of our new exclusive-to-Apress reader with unparalleled search functions
- Option to download any eBook for just $4.99 for a limited time

www.apress.com | @apress

Want more info? Check out www.apress.com/subscription

Apress ACCESS

The articles contained in this magazine are released under the Creative Commons Attribution-Share Alike 3.0 Unported license. This means you can adapt, copy, distribute and transmit the articles but only under the following conditions: you must attribute the work to the original author in some way (at least a name, email or URL) and to this magazine by name ('Full Circle Magazine') and the URL www.fullcirclemagazine.org (but not attribute the article(s) in any way that suggests that they endorse you or your use of the work). If you alter, transform, or build upon this work, you must distribute the resulting work under the same, similar or a compatible license. Full Circle magazine is entirely independent of Canonical, the sponsor of the Ubuntu projects, and the views and opinions in the magazine should in no way be assumed to have Canonical endorsement.

Please note: articles in this magazine are provided with absolutely no warranty whatsoever; neither the contributors nor Full Circle Magazine accept any responsibility or liability for loss or damage resulting from readers choosing to apply this content to theirs or others computers and equipment.

Welcome to another issue of Full Circle.

No Python this month I'm afraid, so I'm filling in with an article on how to get Android apps running in your *buntu distro using the Arc Welder extension for Chrome. It works surprisingly well, but only in Chrome browser. We still have LibreOffice, another LaTeX article and more JavaScript.

If you've got WiFi problems in your abode and would even settle for a network point then have a read of my Linux Labs article where I explain how I set up a remote IP camera using a pair of powerpoint plugs. I bought them thinking they'd never work properly, or not be sufficient, but they're surprisingly good!

Long time readers will know that I'm a KDE man. Sure, I've fell off the wagon a few times (*cough* 4.0 *cough*), but I've always returned when it's at its best. Well, the KDE team have been busy updating the Plasma back end of KDE to version 5. It really is a thing of beauty. Alan Ward has sent in a review of Plasma 5 for you to read through. If you haven't tried KDE (or just haven't tried it in a while) please do. It's criminally under rated in my opinion. Plasma 5 is now the standard in Kubuntu as of 15.04, so you've no excuse!

I'm still thinking of things to do for FCM#100. If you've any suggestions or, better still, articles that are a bit different from the norm then send them to: ronnie@fullcirclemagazine.org.

All the best, and keep in touch!

Ronnie
ronnie@fullcirclemagazine.org

News

Spam-blasting malware infects thousands of Linux and FreeBSD servers

Several thousand computers running the Linux and FreeBSD operating systems have been infected over the past seven months with sophisticated malware that surreptitiously makes them part of a renegade network blasting the Internet with spam, researchers said Wednesday. The malware likely infected many more machines during the five years it's known to have existed.

Most of the machines infected by the so-called Mumblehard malware are believed to run websites, according to the 23-page report issued by researchers from antivirus provider Eset. During the seven months that they monitored one of its command and control channels, 8,867 unique IP addresses connected to it, with 3,000 of them joining in the past three weeks. The discovery is reminiscent of Windigo, a separate spam

botnet made up of 10,000 Linux servers that Eset discovered 14 months ago.

Source: http://arstechnica.com/security/2015/04/30/spam-blasting-malware-infects-thousands-of-linux-and-freebsd-servers/
Submitted by: **Arnfried Walbrecht**

Apple, Linux devices to be decoded at new CBI lab

The Indian Central Bureau of Investigation has got a new specialized forensic lab to decipher and recover data from Apple devices seized from suspects during investigation of cases. The new lab, inaugurated at the CBI academy in Ghaziabad, will be fully equipped with the latest workstations and software to decode the digital information stored in Apple devices, said sources.

So far, the agency had limitations in deciphering Apple and Linux-based devices, which are becoming extremely popular nowadays, with the training of forensic experts emphasizing Windows-based software.

The specialized lab is focused on extracting information from Apple devices, using forensic software, from devices such as the iMac, MacBook Pro, iPad, iPhone and iPods as well as from Linux devices which are like Android-based devices. The laboratory has been given forensic tools for cloning, imaging, password recovery, forensic analysis and internet artifact recovery to enable speedy probes, said an official.

Source: http://timesofindia.indiatimes.com/tech/tech-news/Apple-Linux-devices-to-be-decoded-at-new-CBI-lab/articleshow/47127838.cms
Submitted by: **Arnfried Walbrecht**

HP's Ubuntu 14.04 laptops offer workstation-on-a-budget performance

It's a small but symbolic boost but the Linux section of the UK online website eBuyer is suddenly looking more populated with the announcement last week by HP of three laptops pre-loaded with last year's Ubuntu 14.04 (Trusty Tahr) LTS.

They're not the first Linux or Ubuntu laptops to be sold on the site by a long shot; an entry-level HP appeared on the site over a year ago and that's before factoring in the Thinkpads already on offer from Lenovo. But while pre-loaded Ubuntu has been around for ages, these machines appear to be aimed at a more mainstream buyer than the well-established developer niche.

Anyone can, in principle, load Ubuntu on the similar hardware, but buying pre-configured hardware has important advantages. Everything is guaranteed to work without issue, and then there's the matter of support which is essential for smaller organisations not inclined to tend their own systems should something go wrong. This is the grown-up way to do things.

All three of HP's efforts are based on AMD's under-rated quad-core A4/8/10 processors coupled to 4GB or 8GB of RAM and 1TB hard drives, on a traditional 15.4 inch laptop platform. They are cheap by PC standards, starting at £199, £249, and even the £300 asking price for the top-of-the-range HP455 represents outstanding value for money.

Source: http://www.techworld.com/blog/war-on-error/hps-new-ubuntu-1404-laptops-offer-workstation-on-budget-performance-3610629/
Submitted by: **Arnfried Walbrecht**

Ubuntu phone that works as a desktop PC coming in 2015

Microsoft recently announced that upcoming Windows Phone devices would be able to operate like desktop computers when you connect an external display, mouse and keyboard. That doesn't just mean that you can run smartphone apps blown up for a big screen: you'll be able to use Universal Windows apps which will change their look and feel for large screens and the operating system will look different on TVs too.

But Microsoft wasn't the first company with this idea: Canonical's been working on bringing similar features to Ubuntu for years. A few years ago the company tried to raise $32 million through crowdfunding to build an Ubuntu smartphone that could work as a desktop… but the company didn't meet its goals.

Now Ubuntu founder Mark Shuttleworth says someone else is building one, and it's set to launch in 2015.

Source: https://liliputing.com/2015/05/ubuntu-phone-that-works-as-a-desktop-pc-coming-in-2015.html
Submitted by: **Arnfried Walbrecht**

---

8 Linux Security Improvements In 8 Years

In 2007, Andrew Morton, a no-nonsense colleague of Linus Torvalds known as the "colonel of the kernel," called for developers to spend time removing defects and vulnerabilities. "I would like to see people spend more time fixing bugs and less time on new features. That's my personal opinion," he said in an interview at the time.

So how's that going? Since Morton issued his call, Linux has added several million lines of code and many thousands of patches and new features. The Linux kernel

development process has shown marked improvement on the security front. It was as good as, or better than, most commercial code when Morton issued his 2007 challenge. As InformationWeek checked into its defect-fixing record, it was surprising how many gains have been made in the last three years.

Linux is better than most commercial code. For example, where one defect per 1,000 lines of code is considered quality, Linux in July 2014 had .55 defects per 1,000 lines. Linux also is better than most other open source projects. That didn't happen overnight, and it didn't happen without changes to the kernel process. What has happened with Linux should serve as a standard by which other projects are measured. As concern grows about the security and maintainability of open source code in the Internet's infrastructure, there may be lessons to be learned from Linux's example.

Source: http://www.informationweek.com/software/enterprise-applications/8-linux-security-improvements-in-8-years/d/d-id/1320294
Submitted by: **Arnfried Walbrecht**

Twistlock Launches To Solve Linux Container Security Problems

As the idea of containers gains momentum, there are a couple of problems that increasingly need to be solved – networking, storage, and security being the key three. Twistlock aims to solve the last of those and be part of unlocking far-broader container adoption.

Containers are, of course, a Linux concept that allows the running of multiple isolated Linux systems on a single control host. With Linux containers (instead of creating a full virtual environment) an operating system is shared

across the various containers while running resources are offered to each container in isolation. Linux containers have existed for a long time, but Docker reinvigorated the notion and brought it to a wider audience.

As Docker has made container usage more prevalent, it has, however, also highlighted some issues with Linux containers that make even broader adoption difficult – storage, networking, and security being the three most regularly cited examples. Indeed, much of the justification for vendors suggesting that containers should still be run within a virtual machine relate to the security issue.

Source: http://www.forbes.com/sites/benkepes/2015/05/07/twistlock-launches-to-solve-linux-container-security-problems/
Submitted by: **Arnfried Walbrecht**

Microsoft's new secure boot strategy will suit Linux firms

Microsoft made its intentions known during its WinHEC conference in Shenzhen, China, in March, when it announced that, in the case of hardware that was installed with Windows 10, it would be leaving the choice of having a means to turn off secure boot up to the vendor.

When secure boot was introduced by Microsoft, along with Windows 8, ostensibly as a means to improve security, it mandated that OEMs had to provide a means for turning it off on the x86 platform. It could not do otherwise as it has in the past been convicted of monopolistic trade practices.

Secure boot is a part of the specification for the Unified Extensible Firmware Interface (UEFI), the replacement for the motherboard firmware or BIOS.

The three main Linux companies Red Hat, SUSE, and Canonical — the last-named being the parent firm of the Ubuntu distribution — have each devised ways to support secure boot. While some other distributions also do so, using the same code as that used by these three, many do not.

Thus, if it was impossible to turn off secure boot on a PC, and one wanted to install Linux on it, then the only option would be to use a distribution that supported secure boot.

Source: http://www.itwire.com/opinion-and-analysis/open-sauce/67959-microsofts-new-secure-boot-strategy-will-suit-linux-firms
Submitted by: **Arnfried Walbrecht**

---

10 of the best Linux distros for privacy fiends and security buffs

Linux distributions can be separated into various categories based on use case and the intended target group. Server, education, games and multimedia are some of the most popular categories of Linux distros.

For security conscious users, however, there's a growing niche of distros aimed at protecting your privacy. These distros help ensure you don't leave a digital footprint as you go about navigating the web.

However, for the truly paranoid, privacy distros are only one part of the equation - and the greater part of that equation involves penetration testing distros. These are distros designed for analysing and evaluating network and system security. These efforts feature a vast array of forensic tools to help you test your configured systems for potential weaknesses.

In this article, we've highlighted 10 of the best privacy

and pen testing distros.

Source: http://www.in.techradar.com/news/software/security-software/10-of-the-best-Linux-distros-for-privacy-fiends-and-security-buffs/articleshow/47211893.cms
Submitted by: **Arnfried Walbrecht**

Will $9 Linux CHIP Replace Raspberry Pi?

Technology startup Next Thing is preparing to unleash a credit card-sized personal computer -- dubbed CHIP -- on the world, turning to the crowdfunding site Kickstarter to drum up financial support.

The project, which has already received close to $700,000 out of a $50,000 goal with 26 days to go, allows users to work in LibreOffice, a free and open source office suite developed by The Document Foundation that allows users to save documents to CHIP's onboard storage.

The basic CHIP costs just $9, with a battery, VGA and HDMI adapter, and mobile casing, which can eventually push the price up to just under $50.

CHIP is expected to ship in December of this year, with the full package with accessories available in May of next year.

Source:http://www.informationweek.com/it-life/will-$9-linux-chip-replace-raspberry-pi/d/d-id/1320359
Submitted by: **Arnfried Walbrecht**

Stealthy Linux GPU malware can also hide in Windows PCs, maybe Macs

A team of anonymous developers who recently created a Linux rootkit that runs on graphics cards has released a

new proof-of-concept malware program that does the same on Windows. A Mac OS X implementation is also in the works.

The developers are trying to raise awareness that malware can infect GPUs and that the security industry is not ready for it. Their goal isn't to tip off malicious hackers, but the source code they released, while incomplete and buggy by design, could potentially be built upon and used for illegal purposes.

The problem the developers are trying to highlight lies not with the operating systems, such as Windows or Linux, nor with the GPU (graphics processor unit) vendors, but rather with existing security tools, which aren't designed to scan the random access memory (RAM) used by GPUs for malware code.

Source: http://www.pcworld.com/article/2921092/gpu-malware-can-also-affect-windows-pcs-possibly-macs.html
Submitted by: **Arnfried Walbrecht**

System76 Meerkat: The Perfect Mini PC for Multimedia or Desktop

Every so often I get the pleasure of writing a review where the biggest challenge is finding something, anything, to nitpick in order to ensure I don't sound like I've been bought off by the company. Such is the case with the Meerkat, by System76. A small device that, at first blush, one might think is a toy. I can assure you, this 4.5" by 4.5" device performs with the power of a machine three or four times its size.

The Meerkat greeted me with Ubuntu 15.04. I immediately set about tweaking the device to fit my tastes and needs. The first step: installing my usual lineup of

software. As soon as I began the process, I fell in love with the device. The Meerkat is fast … much faster than a device this size should be.

I moved the Meerkat over to my 60” LED television, sat on the couch, and turned the tiny PC into the ideal multimedia device. Netflix? No problem. Spotify? I think I will. Youtube? Yes please. The Meerkat handled anything and everything I could give it without a single hiccup. Even rendering a video in OpenShot was blazingly fast. And because the connection to the TV was handled via mini HDMI to HDMI, there was no need to run a second cable for sound. It’s all there and it’s all good.

Source: http://www.linux.com/news/hardware/desktops/829053-system76-meerkat-the-perfect-mini-pc-for-multimedia-or-desktop
Submitted by: **Jack Wallen**

Ubuntu jumps into Internet of Things with Acer, GE, and Microsoft

Yes, you read the headline correctly. Microsoft and Canonical are partnering up on IoT.

Old enemies are becoming new allies as technology shifts from the PC/desktop model to first mobile computing and now the Internet of Things (IoT). Canonical, Ubuntu Linux's parent company, is partnering with Acer, DataArt, and Microsoft.

That wasn't a typo. Canonical and Microsoft, which were already working together on bringing Canonical's Juju DevOps tools to Windows and bringing Windows Server to OpenStack, are working with DataArt on an IoT industrial predictive maintenance solution. It will combine the three companies' IoT, cloud, big data,

machine learning, and Docker efforts. To integrate all of this they'll be using "Snappy" Ubuntu apps, DeviceHive, and Juju Charms. Microsoft will also use an Azure service to manage and capture machine data.

Source: http://www.zdnet.com/article/ubuntu-jumps-into-internet-of-things-with-acer-ge-and-microsoft/#ftag=RSSbaffb68
Submitted by: **Steven J. Vaughan-Nichols**

Ubuntu 15.04 hands-on: One giant leap for developers and the cloud, but one small step for the desktop

There are lots of new goodies in Ubuntu 15.04, Vivid Vervet, but most of them are for cloud administrators and DevOps. Ordinary Ubuntu PC users will find only a slightly better desktop experience.

In this release, the distribution boasts a new light-weight snappy Ubuntu Core version for devices, micro-servers, and containers. It also includes updated developer tools and the latest frameworks, languages, databases and packages. This cloud brand of Ubuntu also comes with superior Docker support, Canonical's own new container-based hypervisor, LXD, and built-in support for the Chef DevOps program.

The rest of the Ubuntu desktop's interface, powered by Unity 7.3, looks and acts the same. So, for example, you can use the HUD to search for programs and files no matter where they may be located on your system.

Source: http://www.zdnet.com/article/hands-on-with-ubuntu-15-04/
Submitted by: **Steven J. Vaughan-Nichols**

Ubuntu LXD: Not a Docker replacement, a Docker

enhancement

Sometimes it seems that Canonical, Ubuntu's parent company, can't win for losing. Often accused of trying to force other open-source groups to follow their lead by keeping projects internal until they feel it's ready to be shared with others, when Ubuntu announced its intention to build LXD, a hypervisor for containers, at the OpenStack Summit, the company was immediately accused of announcing vaporware (!), of shoving LXD down other programmers' throats, and of trying to replace Docker.

On the technical side, LXD, pronounced Lex-Dee, is an expansion of LXC, the Linux container technology behind Docker. Specifically, according to Stéphane Graber, an Ubuntu project engineer, LXD is a "daemon exporting an authenticated representational state transfer application programming interface (REST API) both locally over a unix socket and over the network using https. There are then two clients for this daemon, one is an OpenStack plugin, the other a standalone command line tool."

Source: http://www.zdnet.com/article/ubuntu-lxd-not-a-docker-replacement-a-docker-enhancement/
Submitted by: **Steven J. Vaughan-Nichols**

Ubuntu Core Drone Is the First Drone That Has Apps

The world is changing, and it looks like everything will soon be powered by operating systems and apps, and that includes drones, as unlikely as it might sound.

The Erle-Copter was presented a couple of weeks ago, but now its makers have returned with more details about this amazing piece of technology. This is not the average drone you can get at your mall and that can do basically

three things, fly, crash into walls, and break. It's what you might call a smart drone, and it's powered by Ubuntu core.

To make things even more interesting, it's a drone that's powered by Ubuntu Core and apps, which is quite weird if you think about it. On the other hand, the makers of this drone do need to program it and run various "behaviors" on it. The Ubuntu OS was the perfect choice.

Source: http://news.softpedia.com/news/Ubuntu-Core-Drone-Is-the-First-Drone-That-Has-Apps-480826.shtml
Submitted by: **Silviu Stahie**

Open source has to be more than Linux

While the notion of free software has lasted since the days Richard Stallman was sleeping under his desk at MIT, the full thrust of collaboratively and openly licensed software really took off with the advent of Linux.

Linux took a principle and filled in an important technology gap that inspired the filling of a thousand other gaps too. This led to the rise of the venerable Linux distribution, as myriad as consumer-grade platforms such as Ubuntu and Fedora, to server-grade such as CentOS and Debian, and down to the downright weird such as RebeccaBlackOS.

For those of us born in the brine of Linux, openness and a commitment to living and breathing openness have always been common social components. It is not uncommon in the Linux world to use a Linux distribution and entirely open source applications with a few exceptions, such as Skype and Steam.

Despite some rather remarkable projects, desktop Linux has always struggled to get a hold of the market, hovering

at around 1.5% of overall market share. While the server, cloud, and infrastructure siblings to the Linux desktop have gone on to dominate, the desktop has been lagging behind, despite passionate and high-quality efforts from projects such GNOME, KDE, Elementary, and many others.

Source: http://opensource.com/life/15/5/open-source-has-be-more-linux
Submitted by: **Arnfried Walbrecht**

Snappy Ubuntu Linux Now Used in Networking, Refrigerators

Canonical, the lead commercial sponsor behind the open-source Ubuntu Linux operating system, today announced an expansion of its push to embed Linux in everything from phones to refrigerators—and now network switches. The Snappy Ubuntu Core Linux operating system, a minimal version of Ubuntu Linux that provides an improved updating and security model, is designed for embedded devices and the Internet of things (IoT). Snappy was first announced in December 2014 and became generally available with the Ubuntu 15.04 Vivid Vervet release on April 23. To date, Canonical has positioned Snappy as a platform for the cloud and mobile devices and is now bringing it to Open Compute Platform (OCP) network switches. The OCP networking effort kicked off in 2013 as a way to enable an organization to choose the operating system they want to run on an open networking switch. "Switches to us are important and very strategic," Dustin Kirkland, product manager at Canonical, told eWEEK. "There is switch at the top of every rack in every data center."

Source: http://www.eweek.com/networking/snappy-ubuntu-linux-now-used-in-networking-refrigerators.html

Submitted by: **Arnfried Walbrecht**

Solar Sail Spacecraft Runs Linux and Uses SSH, Says Bill Nye

The idea of solar sails was first introduced in popular culture by none other than Carl Sagan, more than 40 years ago. This particular technology was not a priority for scientists in the past decades, with very few exceptions, but The Planetary Society and Bill Nye want to change that by launching a small spacecraft called CubeSat that will be powered by light.

The principle behind the technology is quite simple. The light particles, the photons, don't have any mass, but they do have a lot of energy. The spacecraft deploys a large solar sail, which is bombarded with light. The acceleration is very small, but constant, which means that a spacecraft can achieve great speed in a relatively short amount of time. Even if this particular piece of information is not available on the Kickstarter page, it was revealed in an AMA session on Reddit, which was held by Bill Nye himself. To make things even more interesting, he also said that it would be possible to SSH into the spacecraft.

"The software is Linux-based. I just recently learned that, providing we have a stable link, we can actually SSH into the spacecraft, which I find very cool. The control sequences are automated. There are sun sensors that locate the sun and tack based on that," said Bill Nye on Reddit.

Source: http://news.softpedia.com/news/Solar-Sail-Spacecraft-Runs-Linux-and-Uses-SSH-Says-Bill-Bye-481089.shtml
Submitted by: **Arnfried Walbrecht**

Tails 1.4 polishes up the privacy-obsessed Linux OS trusted by Edward Snowden

Tails, a privacy and anonymity-focused Linux distribution most famously used by Edward Snowden, just released version 1.4.

This Debian-based system is designed to preserve your privacy and anonymity online, providing better protection than just using the Tor browser alone on a typical operating system. How effective is this concealment-centric operating system's tools? Well, in 2012, vulnerabilities for Tails topped the NSA's most-wanted list alongside Tor and TrueCrypt.

Let's dig into Tails' basic capabilities, as well as the new changes.

Tails stands for "The Amnesiac Incognito Live System," and it's designed to be booted and run entirely from a disc, USB drive, or SD card. This ensures no traces of your activity are written to your PC's hard drive. It also means any malware or other surveillance software running on a computer's normal operating system—Windows, for example—won't be involved with the Tails session.

Source: http://www.pcworld.com/article/2923013/tails-14-polishes-up-the-privacy-obsessed-linux-os-trusted-by-edward-snowden.html
Submitted by: **Arnfried Walbrecht**

Meizu Will Sell Ubuntu MX4 Across Europe Soon, Says Canonical

However, the sad news for many of us who wanted to purchase the Ubuntu-powered device from Meizu is that the Ubuntu MX4 smartphone is sold only in China for the

time being, despite the fact that Meizu said a few months ago that it would sell it internationally.

According to Canonical, the Meizu MX4 Ubuntu Edition smartphone will come to the European market soon, but today's announcement says nothing about the device being sold in the US or another continent anytime soon. Also, the phone appears to be targeted at Ubuntu Touch developers at the moment. "In the first of a series of launches, Meizu has announced that the Ubuntu MX4, which will be sold across Europe soon, is immediately available to developers in China," says Canonical.

Source: http://linux.softpedia.com/blog/Meizu-Will-Sell-Ubuntu-MX4-Across-Europe-Soon-Says-Canonical-481404.shtml
Submitted by: **Arnfried Walbrecht**

Serious Red Hat Linux Bug Affects Haswell-based Servers

A recent post by Gil Tene raises the importance of an important, little known patch to Linux kernels that should be reviewed by all users and administrators of Linux systems, especially those who utilize Haswell processors. Tene reports that in particular users of Red Hat-based distributions (including CentOS 6.6 and Scientific Linux 6.6) should apply the patch as soon as possible. Even if your instance of Linux is running in a VM, that VM is most likely hosted on a Haswell machine if is on the popular cloud providers (Azure / Amazon /etc) and would benefit from the patch.

Tene goes on to explain how the flawed code performed (boils down to a switch block missing a default case). The big reason for the problem today is that while the code in question was fixed in January 2014, the flaw was

backported into the Red Hat 6.6 family around October 2014. Other systems including (SLES, Ubuntu, Debian, etc) are also probably affected.

The fix for those systems is only now being distributed and it could be overlooked. Red Hat users should look for RHEL 6.6.z or newer. A key point made by Tene is that the fix has been unevenly distributed as different distributions make specific choices on what goes into their kernel.

Source: http://www.infoq.com/news/2015/05/redhat-futex
Submitted by: **Arnfried Walbrecht**

Command & Conquer

by Lucas Westermann

Meteor

Over the past year, I've written a number of C&C articles touching on web development (from JavaScript frameworks, to CSS preprocessors). Surprisingly, I have received a relatively large number of emails indicating an interest in this topic. Due to this interest, I wanted to spend this month covering Meteor, which is a relatively unique JavaScript platform, and has also been the focus of my own learning for the last 5 or so weeks.

What is so unique about Meteor?

Meteor is unique in the fact that it is both a JavaScript server (like NodeJS), but also ties into the client side, so the application you write exists on both the server and the client. This brings with it an array of performance enhancements, and some security concerns (which can generally be avoided by following Meteor's best practices).

It's also deeply integrated with Cordova and Phonegap, meaning you can easily create native Android and iOS apps from the finished web app. The nature of the platform is to be reactive (changes are seen instantly, without a refresh - similar to AngularJS), and offers plenty of useful features. To find out more, check out their website (link in Further Reading).

What about the downsides?

As both the server-side and client-side are JavaScript, the page will render absolutely nothing if the browser has

JavaScript disabled. You could output some sort of notice using a set of <noscript> tags in the html's <head> tags, but if ostracizing anyone who has JavaScript disabled doesn't work for you, Meteor will not be for you (at least, not without various hacks, if it's even possible). That's the benefit to something like Angular - depending on how you use it, visitors who don't allow JavaScript will at least see the majority of the normal HTML, even if it's not very pretty. In this day and age, the question is simply how many visitors will have disabled JavaScript. In the matter of SEO - Google definitely updated their bots for JavaScript pages, but I'm unsure of Yahoo or Bing. Checking this before developing might be useful.

That being said, there is a plugin called Spiderable that uses PhantomJS to create static HTML, and send it to spiders (search engine crawlers). Perhaps this approach could be adapted for visitors with javascript disabled, though it does not appear to have been attempted anywhere.

When should I use Meteor?

Technically, you can use Meteor at any time. However, most of the plugins and sites I've seen that use Meteor are geared at heavy interactive sites or complete web applications. If you're planning on creating a small static or semi-static web page, you may want to carefully weigh the pros and cons. That being said, Meteor's own page is naturally created in Meteor.

Depending also on what tools you've currently been using - Ruby on Rails, or Node.JS, switching to Meteor may be an a welcome trade-off, as development with Meteor is extremely quick.

I'm intrigued - where can I learn more?

The Meteor homepage has a nice tutorial for getting

started. There are also various tutorials on YouTube and the web. If you're looking for books, I can recommend Discover Meteor (by Tom Coleman and Sacha Greif), whose authors are both active Meteor developers.

Hopefully this article will be helpful to some - especially those who love to tinker with new frameworks. If you have any questions, issues, or requests, please let me know at lswest34+fcm@gmail.com. As always, I'm also open to requests for future articles.

Further Reading

https://www.meteor.com/ - Meteor Homepage
http://www.telescopeapp.org/ - A social framework built in Meteor

LibreOffice Part 49, Macro Security

Elmer Perry

In next month's article, I will discuss the use of macros as functions, but before I do, I must discuss macro security. While macros are a powerful tool when used properly, macros can contain code that is harmful to the data on your computer. Through the years, documents with macros have been the transfer method for many computer viruses. With a little caution and a few settings, you can minimize the chances of your computer getting a virus from document macros.

Security Levels

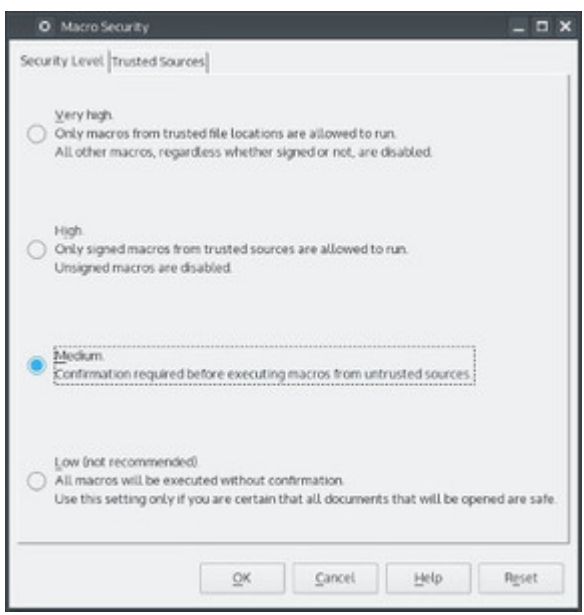
Macro Security
Security Level
Trusted Sources
Very high.
Only macros from trusted file locations are allowed to run.
All other macros, regardless whether signed or not, are disabled.
High.
Only signed macros from trusted sources are allowed to run.
Unsigned macros are disabled.
Medium.
Confirmation required before executing macros from untrusted sources.
Low (not recommended)
All macros will be executed without confirmation.
Use this setting only if you are certain that all documents that will be opened are safe.
OK
Cancel
Help
Reset

Open the macro security settings through the menus: Tools > Options > Security > Macro Security (button). LibreOffice has four different macro security levels that cover security – from no restrictions to highly restrictive.

Each level has its merits. Let's look at each one in detail.

Low: This is the “Off” level. All macros are executed without prompting. They can run without you knowing, and they can damage your files and settings. It's like walking through a snake pit in your bare feet: you will get bitten. The only way this level would ever make sense is on a computer that is completely isolated from the Internet and never opens a file that was not created on that computer. Not very likely.

Medium: This level is the “Are you sure?” level. With medium level protection, the document will run macros if it comes from one of the trusted sources, discussed below, without any prompting. If the document is not from a trusted source, LibreOffice will prompt you about whether to run the macros in the document. You have the choice: Yes or No. I have recommendations later about how to make this decision.

High: This is the “hands off” level. Only signed macros from a trusted source, or macros from trusted file locations, are allowed to run. All others will have their macros disabled. You have no choice; you are never prompted. This is a level where you don't want the end users making the decisions, but you want LibreOffice to make the decision for them.

Very High: This is the “paranoid” level. Only files from trusted file locations can run macros. Again, you are not prompted or given a choice. If the file doesn't come from a trusted file location, the macros are disabled. This is the most locked-down, don't-trust-nobody level there is. If you think the world is out to steal your identity and know all your secrets, you might be right, and this is the level for you. Trusted Sources

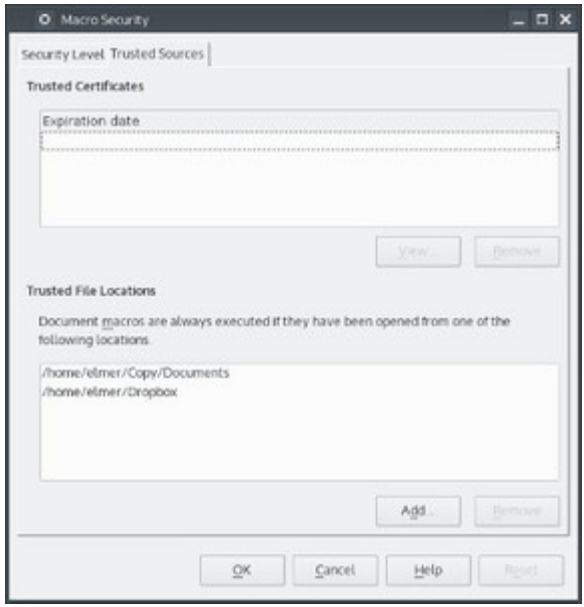

The Trusted Sources tab allows you to identify the trusted sources for your documents.

Certificates: Certificates are used to digitally sign documents. Certificates come from a certification authority. They are usually used on web sites and servers to authenticate the source. The installation of certificates is beyond the scope of this article, but to digitally sign a macro, Tools > Macros > Digital Signature.

File Locations: No matter what level you use – except Low – I recommend you define at least one trusted file location. You need at least one place where you can run proven files with macros without having to OK a prompt. Don't use a location where you usually download files from email or the Internet. Use a location where you collect and save documents you need to keep. Sometimes, two or three locations are helpful. Also, make sure the location isn't too general, like your home folder. Recommendations

Never, ever use Low. No, I mean NEVER. Don't try to argue that you know what you are doing, and you have a firewall, anti-virus, and malware protection. Don't do it! No. No Low level setting. Go there now in your copy of LibreOffice and change it to something else. There. Now, don't you feel safer already? Now, we can calmly discuss the benefits of the other levels.

Medium is my preferred setting. While it protects you, you also get the choice of making that decision for yourself. I recommend this setting for computers on your home network. The setting gives you protection, but does not take away your right to make a choice. If some computers are used by students, you will want to educate them on a good method for making the decision for documents from email and the Internet. Later I will offer you a method that I feel is useful.

High and Very High are best reserved for office networks and computers used by younger students. Use these when you are locking down a computer where you want to limit the power of the end user. I see this in my line of work all the time. The company has the user's computer limited in what they can do to protect the company's investment in the computer and data. Sometimes, this is set according to the experience and necessity of the users. If you are the network administrator, you will need to make this decision for each user.

As I said before, whatever level you decide to use, I recommend at least one trusted file location. In fact, for the Very High setting, you have no choice. Also, with file locations, you can avoid the need to sign the macros in every file with a security certificate.

You Are Your Best Tool

The most important virus protection of all: the mushy

gray matter between your ears. If you were not expecting a document to contain macros, then, maybe, you shouldn't allow the macros to run. Follow your instincts that developed in humans as a protective measure. Be the skeptic! If you have a feeling that something is not right, maybe something is not right.

Take these steps when receiving a new file. Open the file from an untrusted file location, but do not allow the macros to run. Examine the macros in the document. If you have any questions about the macros, asks the person who sent the document. If they cannot give you a satisfactory answer as to why the macros are there, eliminate the document. Chances are, they are not your friend and do not have your best interest in mind (or their computer is infected with a virus). Yes, this may all sound a little paranoid, but better safe than losing time and data. If the document is clean (no macros) or the sender gives you a satisfactory reason for the macros, move the document to a safe file location.

While macros are useful, people also use them for malicious purposes. LibreOffice provides you with four different levels of protection against malicious macros. You should never use the Low level setting, but the Medium level is good for most home computers and networks. High and Very High are good in office environments where you need to control the end users interaction. Always take precautions when dealing with documents coming from other people. Sometimes, people unknowingly pass on documents that contain malicious code.

Drawing with Inkscape – Part 37

By Mark Crutch

Over the past few instalments, I've detailed the workings of the Tiled Clones dialog. As well as producing some interesting, and often kaleidoscopic, visual effects, this dialog can be useful for creating particular arrangements of objects. Because the dialog's units are based on the bounding box of the parent object (rather than using explicitly stated dimensions), this approach isn't terribly useful for positioning objects at specific coordinates or locations, so the next few articles will look at the different tools Inkscape provides for these kinds of manipulations.

It's important to remember, however, that Inkscape is not a “computer aided design” (CAD) program, of the sort used by architects or engineers. If you want to create technical drawings that could be used by builders or manufacturing companies, there are other Open Source programs that might serve you better, such as FreeCAD, OpenSCAD, or QCAD. With each release, Inkscape gains more functionality in this area – such as the new Measure tool in 0.91 – but it's still a long way from a fully fledged CAD application. With that limitation in mind, we'll start with the most obvious form of positioning: snapping to a grid.

Inkscape is rather flexible when it comes to grids. You can have more than one active at a time, and each one can be either rectangular (based on horizontal and vertical lines), or axonometric (vertical lines, plus two sets of angled lines). Usually, however, it's easiest to work with just one grid at a time, most commonly a rectangular one. Launch Inkscape to create a new document, and, if there's no grid visible, try pressing “#” or use View > Grid to enable the display of the default grid. To edit this grid, or to add a

new one, use the File > Document Properties menu, then select the Grids tab.

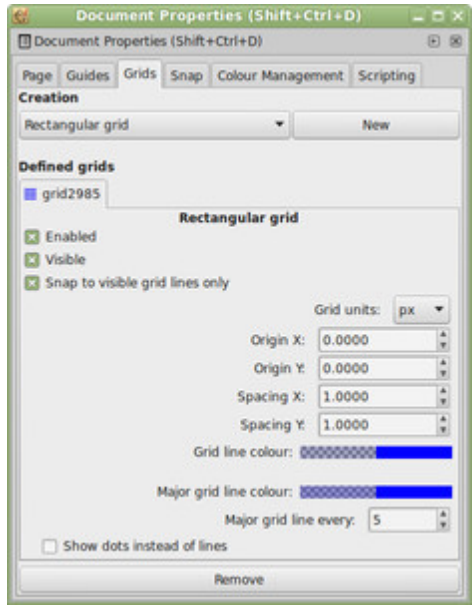

To create a new grid, select either Rectangular or Axonometric from the pop-up menu at the top, then click the New button. Each grid gets its own tab in the “Defined grids” section, with slightly different icons for the two grid types – but the icons are always blue, regardless of the color of your grid lines. They're named automatically, with no easy mechanism to change the name, and no distinction in naming between the two grid types. The use of similar, identically-colored icons and fixed, generic names means that managing numerous grids is less than straightforward, so it's best to stick with just one or two. The tab bar isn't scrollable – the dialog just grows in size if you start to add too many; and there's also a practical limit imposed by the UI.

Each grid has three checkboxes that control snapping and visibility. These are in addition to the View > Grid option (toggled via the “#” key), so, for a grid to be visible, you

need both of the top two checkboxes to be enabled and the View > Grid toggle to be on. Think of the toggle as a global way of showing and hiding all the possible grids at once, with the checkboxes being a way to more finely control what each individual grid contributes.

The first checkbox, “Enabled”, simply switches the entire grid on and off. With it unchecked, the remaining options are all disabled and the grid plays no part in snapping or display. You might use this if you have multiple grids so that only one is switched on at a time – typically when swapping between rectangular and axonometric within the same drawing.

The second checkbox, “Visible”, has only a visual effect. With this unchecked, the grid will not be displayed, but can still play a part in snapping if the grid is enabled via the first checkbox. As it can be difficult and frustrating to snap to a grid you can't see, I strongly recommend always leaving this option checked.

The third checkbox, “Snap to visible grid lines only”, is slightly misleading with its labelling – though the tooltip is clearer. This checkbox concerns grid lines that are automatically suppressed as you zoom out. If you were to zoom out far enough, the grid lines would become so dense that they just appear as a solid colored background. Inkscape avoids this by hiding grid lines that would be drawn too densely, and would have been a hindrance rather than a help. With this option unchecked, you can still snap to these suppressed lines, but I recommend leaving it enabled so that you snap to only visible grid lines. This may require you to zoom in a little to get the precise snapping point you're looking for, but it's usually a good trade-off against the frustration of constantly mis-snapping to invisible grid lines.

The remaining options in the dialog are used to define the

displayed lines, and are fairly self-descriptive. You can position the origin of the grid – 0,0 is good for most drawings – and define the spacing between grid lines in both the x and y directions. As most people usually need a square grid it would be good to have an option to link these values, but it's no great hardship to simply enter the same value into each field. You can set the color, and opacity for the grid lines – it's usually best to keep the opacity low so that they are less dominant on the screen, and less likely to be mistaken for real lines in your drawing. By setting different colors or opacities for major and minor grid lines, you can create a “graph paper” effect. This becomes visible only when you're zoomed in closely enough, and, if you do decide to use this feature, it's best to make the major lines more opaque than the minor ones. Set the “Major grid line every” field to 0 if you want to disable this feature. Finally, the “Show dots instead of lines” checkbox does exactly what it describes, and results in a more lightweight view of the grid. Setting this option, or changing the line colors, can make it easier to work with multiple grids that are all active at the same time, should you need to.

The options for creating an axonometric grid are broadly similar to those for creating a rectangular one. There's only a single entry for spacing, and a couple of new fields for defining the angles of the x and z axes. The default values of 30° are fine for isometric drawings, though you might want to use 45° for an “oblique” projection. Setting either of these too close to 0° or 90° results in rendering problems. Any value outside this range is capped, though that isn't reflected in the displayed number. Unfortunately, there's no “Show dots instead of lines” option for axonometric grids.

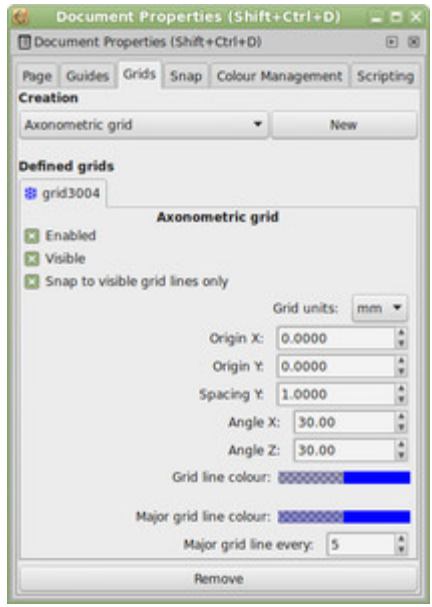

Also missing is the ability to create logarithmic or polar grids. Both are possible via extensions – though these create “real” SVG objects, rather than Inkscape grids, so they can't be toggled with the “#” key, and have to be snapped to via object snapping, not grid snapping. If you do have need of these grid types, however, you can find them under Extensions > Render. In 0.48, you're looking for the Cartesian Grid and Polar Grid entries, whereas in 0.91 they've both been moved into an extra “Grids” submenu. The image below gives an example of grids produced using these extensions.

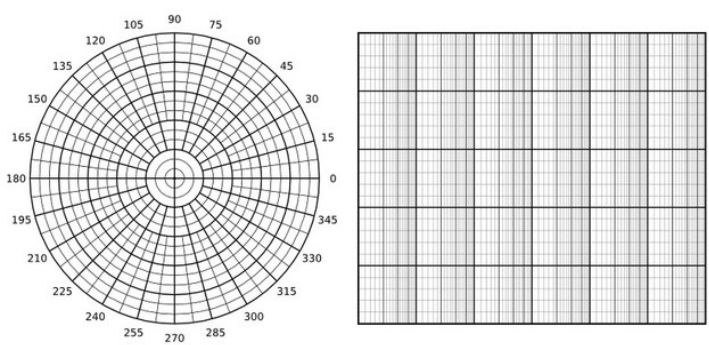

Going back to Inkscape's rectangular and axonometric grids, the setup we've done so far just lets you visually align objects. To really use them as layout tools, you'll also need to enable snapping. As with grids, there's a global toggle for snapping: press the “%” key (Shift-5 on my keyboard) to trigger it, use the View > Snap menu entry, or click on the first icon on the Snap Controls Bar (button 1 in the screenshot).

With snapping enabled, you also need to have either button 2 or 3 (or both) enabled. These dictate what parts of your drawing can be snapped: button 2 enables the snapping of bounding box corners, and kicks in when you resize or move an object using the Select tool; button 3 enables the snapping of nodes and handles, so has an effect when you initially draw an object, or subsequently edit it using the appropriate object-specific tool. Enabling either button makes a few others available that you can use to enable more specific snappings – such as snapping the center point of a bounding box, rather than just its corners – but most of the time I find that one of the main two buttons is good enough.

As well as defining which parts of an object should snap, we also have to tell Inkscape what we want them to snap to. Button 4 enables snapping to the grid. The button after it enables snapping to guidelines (covered in Part 16 of this series) – I usually leave this enabled as it has no effect unless you specifically add some guidelines, in which case you probably want it switched on anyway. The other button in this section enables snapping to the page border, which I find to be less useful for my own projects.

There's one final setting to look at regarding snapping, so it's back to File > Document Properties, but this time

select the Snap tab.

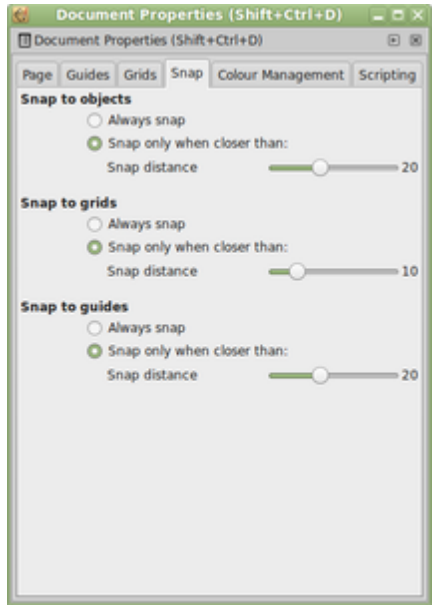

Within this panel, you can set how voracious Inkscape is in its efforts to snap. For each section, the “Always snap” option means exactly that: snap to the nearest snapping point, regardless of how far away it might be. This is useful if you absolutely have to draw to the grid to ensure dimensional accuracy. More generally, however, “Snap only when closer than” offers a good trade-off between accuracy and freedom. With this mode enabled, you can freely place your objects, nodes and handles, unless they get close to a snapping point, in which case they will jump to that location. Exactly how close is set by the “Snap distance” slider. The values are measured in screen pixels, so zooming has an effect on the effective “hit area” in which snapping occurs: zoom in to give yourself more freedom, or zoom out to make snapping more likely.

Within File > Inkscape Preferences (Edit > Preferences in 0.91), there are “Grids” and “Snapping” panels that offer a few other options. The defaults are usually fine for

most users, but if you find yourself heavily using grids or snapping, it might be worth taking a look to see if any of these settings can improve your workflow.

With a grid visible and snapping enabled, you should find it quite easy to create shapes that stick to the grid intersections. When you need to place nodes off the grid, you can either press the “%” to disable snapping (whilst leaving the grid visible), or more easily (on a US or UK keyboard, at least) you can press “#” to turn the grid off entirely, removing it as a snap target. Often, however, there's a requirement to place objects relative to one another, rather than to an absolute grid; that will be the subject of the next instalment...

Mark uses Inkscape to create three webcomics, 'The Greys', 'Monsters, Inked' and 'Elvie', which can all be found at http://www.peppertop.com/

Programming JavaScript – Part 2

By Gergo Bogdan

In the first part, I talked about how JavaScript has evolved over the years, what is the current status of the language, and why it is so widely adopted by developers. In this part, I will focus on JavaScript Objects and functions. Since I received an email from Ray (one of our readers, thank you for contacting me) with some questions, I will address these at the end of the article in the Questions and Solutions section.

JavaScript Objects

In JavaScript, everything is an object, even functions are objects, which, if you are familiar with other programming languages, may seem a little odd. But, don't worry, this gives JavaScript some real power.

JavaScript is a dynamic programming language, which means objects can have different values assigned during runtime, without any problem. For example:

```
var myName = "Greg";
myName = 123;
```

In the first case, the myName stores a string value; in the second case it should store an integer, 123, but JavaScript has a special type, Number (https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/Number) to hold numerical values.

Since JavaScript is a dynamic language, I can extend objects in any way I want. For example, if I create a new object:

```
var myHouse = { nrOfRooms:3, size:100 };
```

It creates an object:

```
Object {nrOfRooms: 3, size: 100}
```

Let's extend this with a couple of properties:

```
myHouse.price = 1500;
myHouse["currency"] = "USD";
```

These result is the following object:

```
Object {nrOfRooms: 3, size: 100, price: 1500, currency: "USD"}
```

As you can see, there are three ways to create custom objects:

1. The first, also called JSON (JavaScript Object Notation), defines objects using the curly braces and specifies the properties and their values separated by a colon.
2. The second option is to use the . (dot) operator, and write the name of the new property and assign a value to it.
3. The third option is to use the index [] operator, which receives a string as a parameter and a value to assign. If you know other programming languages, you can imagine JavaScript objects as a kind of special dictionaries or maps.

Functions

Let's look at the code below:

```
myHouse.getInformation = function() { console.log("The house is " + this.size + "
mm2 big and it has " + this.nrOfRooms + "."); }
```

If you invoke this function as:

```
myHouse.getInformation();
```

it will display : The house is 100 mm2 big and it has 3.

Functions can have parameters:

```
myHouse.isBiggerThan = function(otherHouse) {
return this.size > otherHouse.size;
}
```

We can invoke the function:

```
myHouse.isBiggerThan({size:98});
```

or

```
myHouse.isBiggerThan({size:198});
```

In the first case, it returns true and in the second case false. Notice, I passed in only a simple object with one property, size, so the code can execute correctly, but I could have written this too:

```
myHouse.isBiggerThan({size:198, nrOfRooms:5, price:8500, currency:"EUR"})
```

or

```
var friendsHouse = {size:198, nrOfRooms:5, price:8500, currency:"EUR"};
myHouse.isBiggerThan(friendsHouse);
```

The result will be the same.

Functions may or may not have return values. In the case of the isBiggerThan() function, I have not specified a return type, nor that it will have a return value, but I could easily return a boolean value (true or false), JavaScript permits this.

As an exercise, you can create other objects which simulate real life objects, like a forest which has a function called plantTrees, and receives a parameter nrOfTrees, and it sums up the number of trees in the forest. Or it may store the different types of animals

which live in the forest; the topic does not really matter, the idea is to get you familiarized with object notation and function creation, we will use this a lot.

Questions and Solutions

Question: Ray asked how we can select some information from a file, especially a SQLite database using JavaScript.

Solution: There are two approaches, two scenarios. The first one is when you have the SQLite database available on the client machine, where the browser is running. In this case you can use the SQL.js (https://github.com/kripken/sql.js/) library to load the SQLite files. Even more, you can do queries and create data schemas using JavaScript. The page on GitHub has good code examples; you can start using those for loading your database and querying data.

The second scenario is when the SQLite database is available only on the web server and the client is accessing the web server through the Internet. In this case, you need to have a server-side component (this can be node.js based) which reads the SQLite database, and, using HTTP requests, sends the data back to the client's browser. In this case, the client side JavaScript is more complex, because it needs to use AJAX calls to load the data. Here is a sample node.js code which connects to a SQLite database, creates a new table if it does not exist, and inserts two new entries in the table. After the insert it queries the table and writes the data to the console:

```
var fs = require("fs");
var file = "rays_data.db";
var sqlite3 = require("sqlite3").verbose();

//
// create the db instance
//
var db = new sqlite3.Database(file);
```

```
db.serialize(function() {

//
// create the table if the database file was missing.
//
if(!fs.existsSync(file)) {
db.run("CREATE TABLE People (firstName varchar(50), lastName varchar(50))");
}

//
// create a parameterized, prepared statement
//
var stmt = db.prepare("INSERT INTO People VALUES (?,?)");

//
// run the statement twice
// each ? in the statement is substituted with the parameter
//
stmt.run("John", "Doe");
stmt.run("Jane", "Doe");

// close the statement
stmt.finalize();

//
// query the database and log the result to the console
//
db.each("SELECT rowid AS id, firstName, lastName FROM People", function(err,
row) {
console.log(row.id + ": " + row.firstName + " " + row.lastName);
});


});
```

To execute the above code, you will need to have node.js (https://nodejs.org/) and npm installed, and install sqlite3 through npm using this command:

```
npm install sqlite3 –save
```

When all is set up, you can run the program using:

```
node sqlite_reader.js
```

Supposing you named your file sqlite_reader.js.

The output should look something like this in the command line (I was executing under Windows, but the

output is the same under Linux too, since node.js runs the same way in both environments):

```
C:\Users\Gergo\Documents\nodesqlite>node sqlite_reader.js
1: John Doe
2: Jane Doe

C:\Users\Gergo\Documents\nodesqlite>dir
Directory of C:\Users\Gergo\Documents\nodesqlite
08/05/2015 21:09 <DIR> node_modules
08/05/2015 21:36 2,048 rays.db
08/05/2015 21:23 953 sqlite_reader.js
2 File(s) 3,001 bytes
```

In future articles, I will present how to create a new HTTP server using node, and how to transfer data from the server side to the client side – everything using JavaScript.

I would be happy to hear from you; what are the topics in JavaScript which you would be interested in. Please feel free and email me your topic ideas, the same way as Ray did. Thanks again, Ray!

Bio: My name is Gergo Bogdan, I am a software engineer, blogger, tech enthusiast, from Budapest, riding the waves of the constantly changing IT ocean ☺, you can check my website http://grelution.com.

Run Android Apps In *buntu

by Ronnie Tucker

Ever fancied having that handy Android app in your Linux distro? Or do you need an Android app on your desktop for a presentation, or for testing? Well, Google's 'ARC Welder' extension is just what you need. Although, if you're using anything but Chrome, you're out of luck. This requires Chrome.

If you haven't already got it, you can get Chrome from: https://www.google.com/chrome/browser/desktop/

First thing is to Install the ARC Welder extension: https://chrome.google.com/webstore/detail/arc-welder/emfinbmielocnlhgmfkkmkngdoccbadn (or search for it at: https://chrome.google.com/webstore/category/extensions)

Click the 'Add to Chrome' button at the top right of the pop-up window.

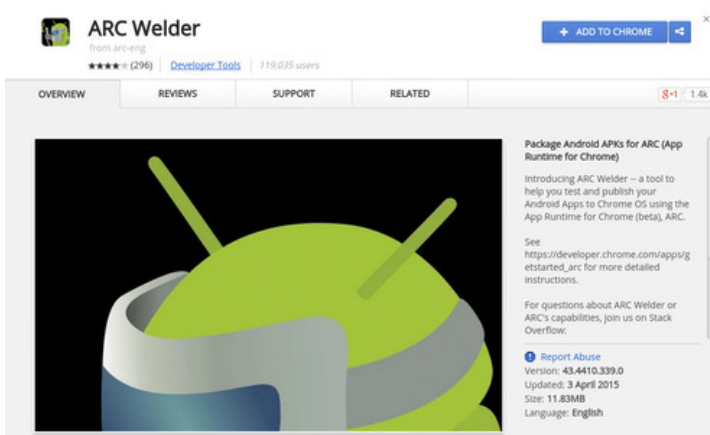
ARC Welder
ADD TO CHROME
Developer Tools
OVERVIEW
REVIEWS
SUPPORT
RELATED
Package Android APKs for ARC (App Runtime for Chrome)
Introducing ARC Welder -- a tool to help you test and publish your Android Apps to Chrome OS using the App Runtime for Chrome (beta), ARC.
See https://developer.chrome.com/apps/getstarted_arc for more detailed instructions.
For questions about ARC Welder or ARC's capabilities, join us on Stack Overflow.
Report Abuse
Version: 43.4410.339.0
Updated: 3 April 2015
Size: 11.83MB
Language: English

You'll get a final confirmation box before it installs. Add

“ARC Welder”? Absolutely!

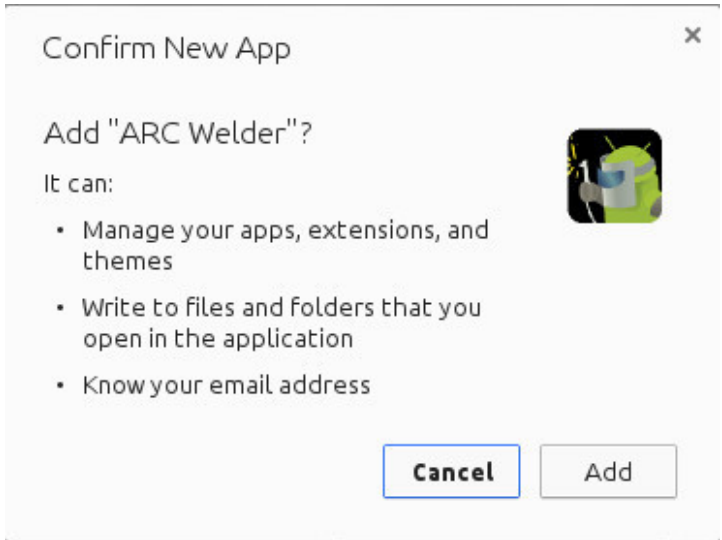
Confirm New App
Add "ARC Welder"?
It can:
• Manage your apps, extensions, and themes
• Write to files and folders that you open in the application
• Know your email address
Cancel
Add

The original blue ‘Add to Chome’ button will change to say ‘Checking…’ for a few seconds. After that’s gone, you’re done. You now have an ARC Welder. One that won’t burn your house down. Yay!

KDE USERS:

In KDE (Kubuntu 14.10 that I was using), if I click ‘App Runtime for Chrome’ shown in my K menu, it pops up with a window of my installed extensions. Instead I had to use the Chrome ‘Show Apps’ button (top left in Chrome), right-click the ARC Welder icon, choose ‘Create Shortcuts’, tick both boxes, and click Create. This now puts an ARC Welder icon in the Chrome Apps folder on the K menu. You need to do this only once.

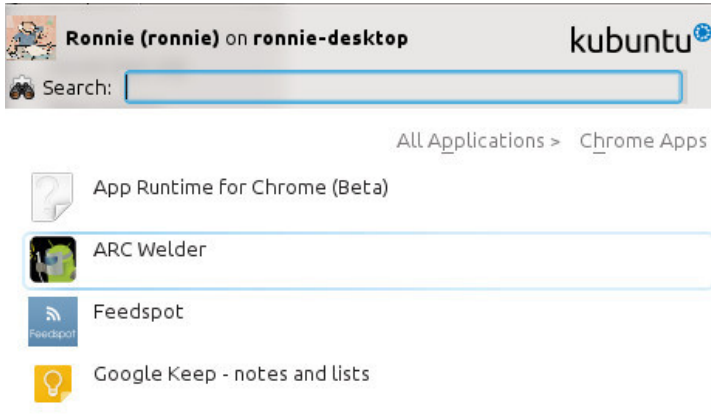

So, let's open ARC Welder. This is either done using your desktop menu, or using the icon in the Chrome apps tab.

ARC Welder
ARC Welder
Welcome
ARC Welder needs to write files to a directory on your filesystem. Please choose a directory.
Choose

Click the yellow Choose button. As the pop-up says, it needs a directory to save stuff in. Give it a directory, and we're off to the races.

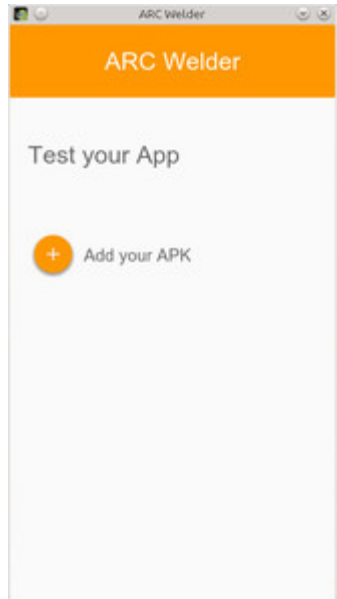

Well, almost off to the races. We need Google Android APK files. First, let's choose an app to test this with. Go to the store (https://play.google.com/store) and choose an app.

I'll use News Republic for this tutorial. Copy the URL of the app. Go to the site: http://apps.evozi.com/apk-downloader/ and paste in the app store URL.

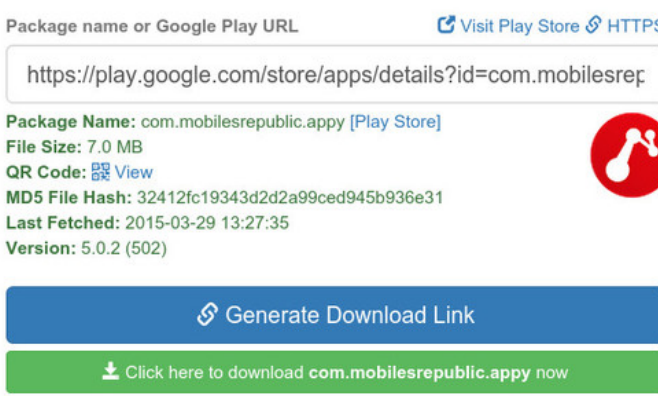

Click the green download button and save the APK file to your hard drive.

With the APK downloaded, it's back to ARC Welder. Click the 'Add your APK' button and choose the downloaded APK file.

ARC Welder will now give you some options on how to display the app. I'll choose landscape, tablet, and give it clipboard access.

Finally! Time to click LAUNCH APP (in yellow at the bottom of the ARC Welder window).

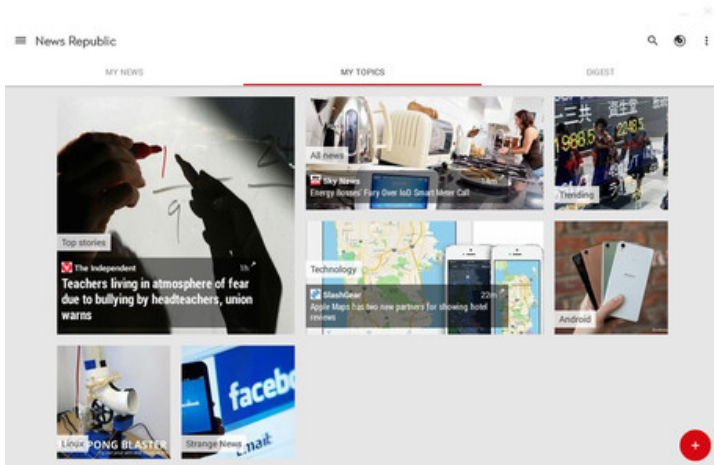

Yay! It worked!

Just remember that you need to click and swipe as you would with the touch-screen version.

Using Latex - Part 3

by John Kerr

In the letters section of Full Circle #96, fellow reader Tiago wrote in to say that TexStudio is an excellent LaTeX text editor. I agree; I also use TexStudio and it will be the LaTeX editor of choice for this month's installment. Gummi and LaTeXzila should not be left out either, and the choice of software may come down to what is the best tool for the job, or for you. Gummi may be the best writing tool for an essay or a quick document that will not require special formatting.

Thank you Tiago for your letter.

The LaTeXzila project is looking for funding to create a real time file viewer and other improvements, see https://wiki.gnome.org/Apps/LaTeXila/donate for more information.

An item to add to your documentation library is a LaTeX cheat sheet. Cheat sheets have many of the commands that you would use. There are standard cheat sheets: http://www.stdout.org/~winston/latex/latexsheet.pdf as well as specialised ones like this for math symbols: http://estudijas.lu.lv/pluginfile.php/14809/mod_page/content/12/instrukcijas/matematika_moodle/LaTeX_Symbols.pdf. A Google search will undoubtedly uncover more.

A cheat sheet is handy to have with a basic text editor or an app like Gummi. However, one of the features of TexStudio (and why Tiago and I like it so much) is that TexStudio has a built-in cheat sheet that can automatically fill in the commands as you type them, or, with a click of the mouse, show you a drop-down menu of

commands and command categories.

Figure 1 shows us how useful the menu can be, as they organize the commands by category and in a natural language. If you know the command, you can still start typing it, and TexStudio will fill in as you type. How many commands you memorize will depend on how often you create documents in LaTeX, but LaTeX is so powerful, you will get help from the menus in TexStudio sooner or later. Not everything is there; for instance drawing a double line is a topic that you will have to search in Google in order to do it.

Figure 2 shows us the Quick Start dialog found under the Wizards dialog for creating the preamble of your document. The menus in TexStudio are extensive and help our workflow when our memory fails us.

Make Your Own Business Cards

I often say that “mountains I can move today, little molehills can take forever”. I just never get around to making business cards for myself, and the fact that I would not hand out 500 cards in a hundred years has not helped the project at all. At a conference I attended last week, I discovered that I needed business cards, and since my hotel had a public access printer, I decided to make some on my laptop using LaTeX.

Find a Template

There is no point in reinventing the wheel, many word processing packages provide templates, and the LaTeX community has provided many as well. A Google search on “Latex business cards” pointed to a number of templates. I downloaded three of them and found that there are many different approaches to creating business cards. What is common is that they all have a section that places the data fields in your card, and a section where

you may enter the information for your card.

A plain business card template that I found[1] has this personal data section:

```
%%%%%%%%%% DEFINE USER SPECIFIC MACROS BELOW %%%%%%%%%
\def\Who     {}
\def\What    {}
\def\Where   {}
\def\Address {}
\def\CityZip {}
\def\Email   {}
\def\TEL     {}
\def\FAX     {}
\urldef{\WEB}\url{}
```

Within the brackets above you will insert your data, save and then compile; the file places the data on the card.

However, using three different editors, and different files, I discovered that, with business cards, another common "feature" is that, after you save and compile, the cards do not appear in the editor's PDF viewer. To see the results of your work you will have to open the PDF file with your computer's PDF viewer.

Try different files, and modify them to make your own unique card.

ShareLaTeX

While searching for templates you may find yourself landing on a site called ShareLaTeX[2]. This site is worth investigating as it provides you with an environment to create LaTeX documents online (or in the cloud). You will be able to store your documents on their servers and edit them from anywhere using their on-line editor. Single user accounts are free, if you want to collaborate with

others there is a charge. It is the collaboration feature that would attract many universities to using the service. It would also be hard for students to lose homework!! Another advantage to any school is with everyone using the same tools on-line, the ability to support LaTeX is easier and they can do it on all computer platforms because the service is web based. However, if everyone created documents using TexStudio it would be easier to support LaTeX as well.

Take a look at ShareLaTeX if not to just view the templates; you do not have to avail of their service to use the templates as you can download them to your computer.

Until next time enjoy exploring LaTeX.

1 http://www.netpurgatory.com/downloads.html#cards
2 https://www.sharelatex.com/

Asteroid Light Curve Simulator

by Dr. Laurent ALDON

The aim of this mounting is to simulate the light curve observed by astronomers over time. In astronomy, a light curve is a graph of light intensity of a celestial object or region, as a function of time. In the case of asteroids, the study of the light curve, together with other observations, can allow us to reconstruct the shape of the object. This project was developed for E. Barres School (Le Cres City) in the framework of the French Science Fair.

Setup of the device

A rough piece of basalt rock was drilled and mounted on the axle of a stepper motor. The surface reflects some light depending on the angular position. Readings on analog input A0 were read using the serial port, to log the steps the motor has taken. A 16x2 LCD display (I2C protocol) was used to show both angle and light measurements.

Figure 1: Arduino, Stepper Motor, Light Dependant Resistor, Basalt rock, LCD 16x2

Results

Young children at school were very impressed to see the piece of basalt starting to rotate. After recording some data, we have plotted the measurements as shown in figure 2.

Code http://pastebin.com/rJfVtWny

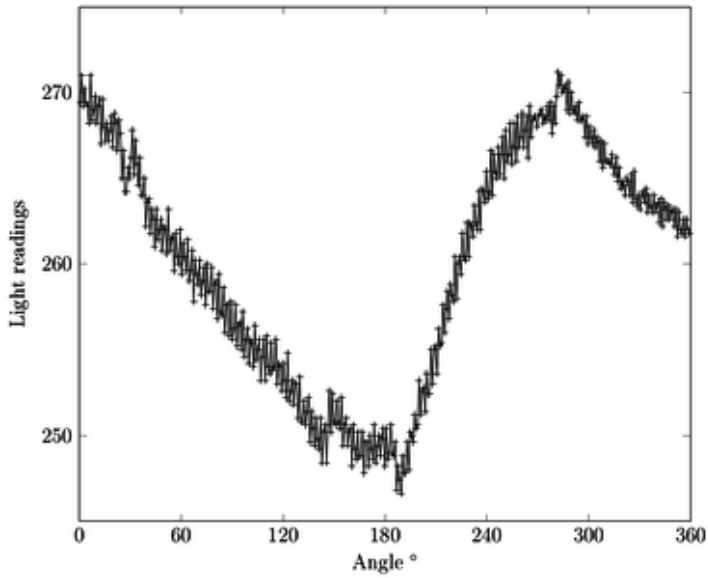

Figure 2 : Measured light curve from LDR and 10k resistor voltage divider.

Arduino code

```
#include <Wire.h>

#include <LiquidCrystal_I2C.h>

// Set the LCD address to 0x27 for a 16 chars and 2 lines display
LiquidCrystal_I2C lcd(0x27, 16, 2);

#include <Stepper.h>

const int stepsPerRevolution = 200;
```

```
// initialize the stepper library on pins 8 through 11:
Stepper myStepper(stepsPerRevolution, 8,9,10,11);
int stepCount = 0;    // number of steps the motor has taken
int light=0;
float angle=0.;
void setup()
{
  Serial.begin(9600);
  lcd.begin();
  lcd.backlight();
  lcd.print("*  Asteroids *");
  delay(2000);
  lcd.clear();
        lcd.print("E. Barres School");
        lcd.setCursor(0,1);
        lcd.print("Falling     2014");
        delay(2000);
}

void loop()
{
  myStepper.step(1);
  light=analogRead(A0);
  angle=(stepCount%2048)/5.625;
  Serial.print(angle);
  Serial.print(" , ");
  Serial.println(light);

lcd.clear();
  lcd.print("Angle : ");
  lcd.print(angle);
  lcd.print("    ");
  lcd.setCursor(0,1);
  lcd.print("Light : ");
  lcd.print(light);
  lcd.print("    ");

stepCount++;
  delay(50);
}
```

SPIDEROAK
.com
Online BACKUP
Secure SYNC
Easy SHARING
Whether you need to access a document you have stored on a remote server, synchronize data between a Mac, Windows or Linux device, share important business documents with your clients, or just rest easy knowing all of your data is safely, securely, and automatically backed up - SpiderOak's free online backup, online sync and online sharing solution can handle all your needs!
SpiderOak offers a different approach to online backup by combining a suite of services into one consolidated tool - free online backup, synchronization, sharing, remote access, and storage. This difference is further measured in our zero-knowledge privacy policy - the first one ever employed in this setting. Our flexible design allows you to handle data from any operating system (Mac, Windows and Linux) or location (external drives, network volumes, USB keys, etc...) using just one centralized account.
Access all your data in one de-duplicated location
Configurable multi-platform synchronization
Preserve all historical versions & deleted files
Share folders instantly in web ShareRooms w / RSS
Retrieve files from any internet-connected device
Comprehensive 'zero-knowledge' data encryption
2 GBs Free / $10 per 100 GBs / Unlimited devices
https://spideroak.com
Download mobile clients for iOS & Android
JOIN SPIDEROAK NOW
Get 2 Free GBs
Get 25% off any SpiderOak package with the code: FullcirclemagFans

Write For Full Circle Magazine

Guidelines

The single rule for an article is that it must somehow be linked to Ubuntu or one of its many derivatives—Kubuntu, Xubuntu, Lubuntu, etc.

Write your article in whichever software you choose—I would recommend LibreOffice. But, **please spell and grammar-check it!**

The Official Full Circle Style Guide can be read at: http://url.fullcirclemagazine.org/75d471

Please read this document before submitting an article. Follow the guidelines and you will have a much better chance of seeing your article in Full Circle.

Writing

There is no word limit for articles, but be advised that long articles may be split across several issues. In your article, please indicate where you would like a particular image to be. Please do not use any formatting in your document.

Images

Images should be no wider than 800 pixels, in JPG format, and use low compression.

When you are ready to submit your article, please email it

to: articles@fullcirclemagazine.org

Non-English Writers

If your native language is not English, don't worry. Write your article, and one of the proofreaders will read it for you and correct any grammatical or spelling errors. Not only are you helping the magazine and the community, but we'll help you with your English!

FOR REVIEWS:

Games/Applications
When reviewing games/applications, please state clearly:
• title of the game
• who makes the game
• is it free, or a paid download?
• where to get it from (give download/homepage URL)
• is it Linux native, or did you use Wine?
• your marks out of five
• a summary – with positive and negative points

Hardware
When reviewing hardware, please state clearly:
• make and model of the hardware
• what category would you put this hardware into?
• any glitches that you had while using the hardware?
• easy to get the hardware working in Linux?
• did you have to use Windows drivers?
• marks out of five
• a summary – with positive and negative points

You don't need to be an expert to write an article – write about the games, applications and hardware that you use every day.

Linux Labs
IP Camera via Powerline

by Ronnie Tucker

My father recently bought chickens and I thought it would be nice to have an Internet camera out there to keep an eye on them. Nice idea, but how can it be done when these Wi-Fi enabled cameras are out of Wi-Fi range? I can't run a giant network cable through the garden.

The answer is powerline adapters. I was highly sceptical of them at first, but I have to admit that they do work, and surprisingly well.

DISCLAIMER: I'll have to be a bit vague in this article as there are many different types of powerline adapters and they all work differently. Same with IP cameras and routers.

Powerline

The idea behind the powerline adapter is this: you buy a pair of power outlet plugs which have network sockets on them and several flashing indicator lights. You plug your first one into the wall and have its network socket wired to your internet router. Your second plug goes wherever you need a network point. In a chicken coop, in my case.

TIP: when pairing the plugs together, do it in the same room as it will save you a lot of legwork.

Most come with encryption built in that you can enable. This is handy if your power source is shared with other residents. This is not something I've tested, but to enable it I think it requires the manufacturer's software, which I found was Windows only and wouldn't run in Wine.

You can add more plugs to your powerline network, but try to keep them all the same type. Otherwise, you're asking for incompatibility and nightmares.

IP CAMERAS

IP (Internet Protocol) cameras come in many shapes, sizes and prices. There are two major manufacturers: Wansview and Foscam.

Yes, they're cheap makes with only 640-pixel wide images, but I don't want an expensive camera in a chicken coop.

The good thing about these cameras is that they pan and tilt. In other words: you can move them by remote control. Most come with infrared LEDs on the front which will allow night vision, too. The back of the camera has an ethernet port, Wi-Fi antenna, power socket, and usually several connections for sound and alarms (if required). All we really need here are the power and ethernet ports.

So, you've got your powerline adapter plugged in where you want it. Plug your network cable from the powerline to the camera, then power on the camera. This is important as I find that powering on the camera first, then plugging in the network cable, doesn't always work for getting an IP address. Give the camera at least one or two minutes before writing it off, as mine takes at least a minute to boot up and go through its up/down, left/right

test sequence.

IP ADDRESS

Most IP cameras come with a sticker on the underside of it. This gives you the default admin password (for initial setup), a dynamic DNS url, and (sometimes) a default IP address. My one didn't come with an IP address on the sticker, so I've no idea what it's IP address is so I'll need to find it. You can, sometimes, get it using the dynamic DNS URL, but they're quite often in China, incredibly slow, and not exactly trustworthy. So, I'd recommend not using it and setting up a faster dynamic URL, but we'll come to that later.

Again, I have to be a bit vague here as all routers are different. I'm using a BT HomeHub 4 (in the UK) so your router will no doubt be different. But, either way, you log in to your router and head to the DHCP table page (if there is one) or to a visual list of attached devices. In there you'll see the device and its IP address.

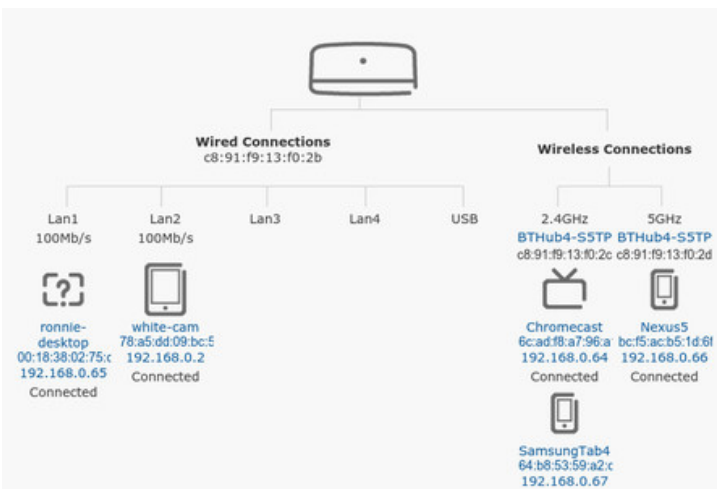

In my case, it was a device I didn't recognise, and I knew that my main powerline was in LAN2 (on my router), so I

knew right away that was the camera. It won't always have an easily recognisable name!

To test things I entered the IP address (192.168.0.2 in this case) into my browser. This should let you log into the camera with the default details which you can, of course, change.

Success!

So, I can access the IP camera over the network. First thing I did, which isn't a necessity, is to head back into my router and tell the router that I want that device to have the same IP address every time. With that saved I know that my camera will always be 192.168.0.2.

By default the IP cameras use port 80 which is the default internet port – that needs to be changed. In the camera settings will be a place to edit the IP address and/or port number. I've changed mine to port 82.

I also changed the password for the admin account, and

added a guest operator account. This means I can give people the guest login to let them view, but not edit, the camera.

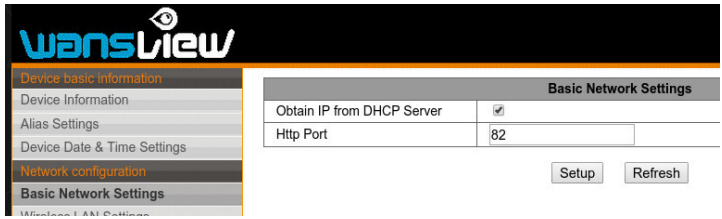

If you plan on having more than one camera online, then each one needs to use a different port number.

So, now that the port is changed, you need to tag that onto the URL. My login now is http://192.168.0.2:82. The colon 82 forces it to use port 82.

Thus far we can log in to the camera on the network, but what about from outside of the network? We need a dynamic DNS for that.

DYNAMIC DNS

If you're like me, then your IP address will change each time to reconnect to the internet. Most routers have a setting for dynamic DNS which will let your router ping its new IP address to a provider who will associate that IP address with an easy-to-remember URL.

There are many different free dynamic DNS providers out there, but I use DTDNS.com as this is what's compatible with my router. You'll need to log into your router and go to the dynamic DNS page to see what it will/won't accept.

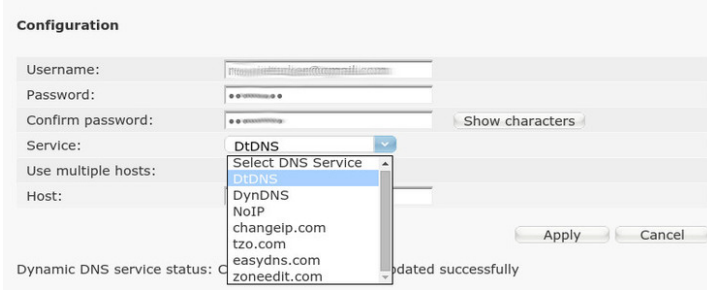

Whichever service you want to use, create an account with them and sign in. Under hostnames you'll choose a name and one of their URLs. I chose, funnily enough, chickens, then one of their URLs.

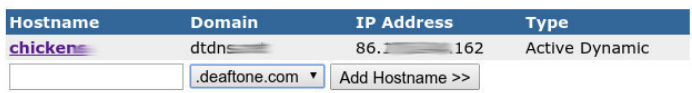

Head back again to the router page for dynamic DNS and give it your username and password for your hostname.

Now, when I reconnect to the internet my router will give DTDNS my new IP address and I can use my URL to log into my camera.

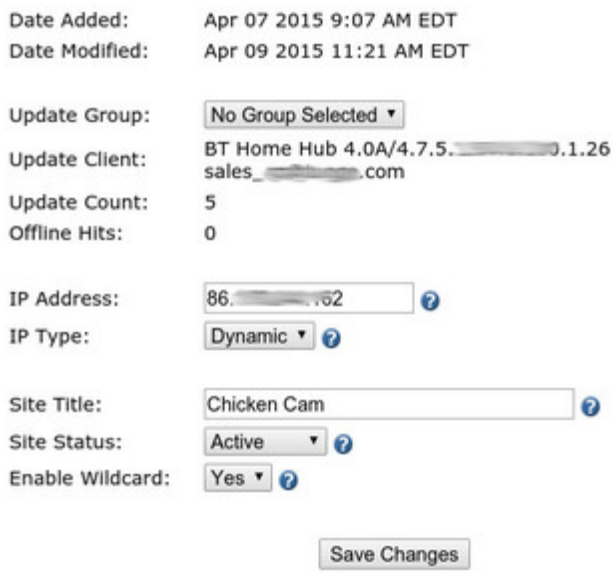
Date Added: Apr 07 2015 9:07 AM EDT
Date Modified: Apr 09 2015 11:21 AM EDT
Update Group: No Group Selected
Update Client: BT Home Hub 4.0A/4.7.5. .1.26
sales_ .com
Update Count: 5
Offline Hits: 0
IP Address: 86. 62
IP Type: Dynamic
Site Title: Chicken Cam
Site Status: Active
Enable Wildcard: Yes
Save Changes

But wait. Typing in the dynamic DNS URL doesn't load the camera!

PORT FORWARDING

When you enter your dynamic DNS URL, you'd expect to log into your camera. But no. Why? Because the URL is using port 80 by default. We need to add a colon and your camera's port number (eg: :82), but first we need to tell the router what to do when we enter that URL.

Somewhere in your router will be a page for port forwarding. It may come under firewall. What we're going to do is first create a label, or name, for each camera. So I add a new application (as my router calls it). For protocol I keep it at 'any,' and I enter 82 for port range, 82 for translate to, and save it. Now I choose my application again (ie: the camera) which is associated with port 82 and link it to the device (the camera).

Yes, it's a bit confusing at first, but the router needs to

know where to send things. When our DNS URL comes in, the router has no idea where to send that port 82 data. We're telling it here to send that port 82 data to the camera.

Now, when I enter my DTDNS URL with the :82 at the end I go straight to my camera.

MOBILE APPS

There are many different apps you can use for viewing IP cameras, but for Android I use tinyCAM Monitor (which is free).

I add a new camera and give it the relevant settings such as make, model, DNS URL, port number, and my guest login.

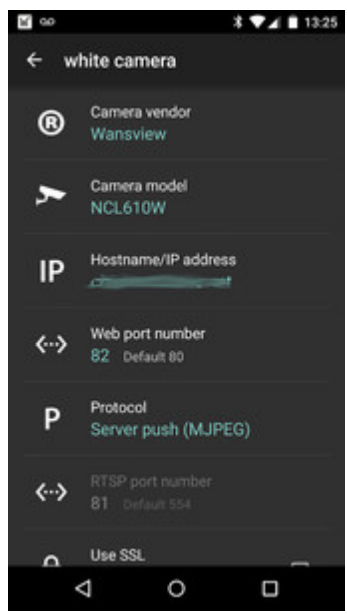

Not in the screenshot shown here, but at the top of that screen is a check to see if the camera is detected, found, and online. If it says it's all OK, then you're good to go.

Go back and choose your camera from the list, sit back, and watch the show.

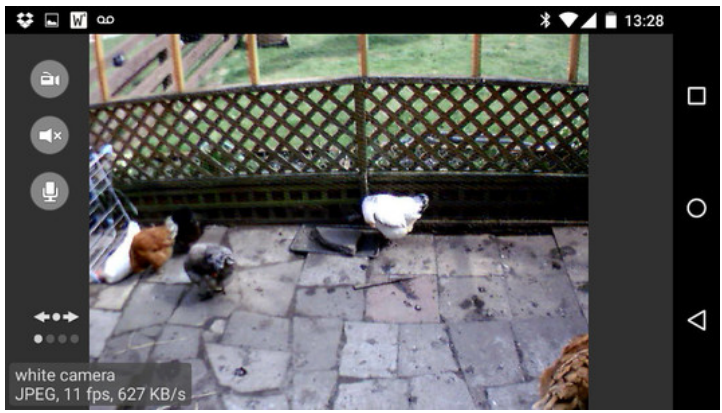

They may be only chickens, but they're far more interesting than most of the current TV shows.

CONCLUSION

Obviously, you can take this further with several cameras throughout your house or property. You can, if you wish, add microphones to your cameras and listen to what your camera is hearing. Some cameras even allow speakers to be added to let you use it as an intercom device of sorts. Even alarms can be wired into some cameras.

The main star of the show is the powerline adapters. Without them there'd be no camera in the chicken coop. A testament to their power is that the coop is quite a distance from the house (ie: long power cable) and the powerline adapter itself is plugged into a four-way extension cable. This is normally a no-no, but it's a necessary evil as there's only one socket in the coop area.

I was a powerline sceptic, but now I'm a believer!

[PS - Footnote from Mike Kennedy (FCM

proofreader): I've used powerline adapters at various customer sites - usually very successfully. However, sometimes they didn't work, or worked poorly. On investigation, the cause was usually a faulty PSU (power supply) in some other device, typically in a very cheap desktop PC, photocopier, printer, etc. So, if the powerline devices are problematic, try switching off most/all other devices, until the troublesome one(s) are identified. Then, perhaps the PSU is actually faulty (though, seemingly still "working"), or perhaps it's of very poor quality - and it should be repaired/replaced].

My Story

by Roland Hunt

After being a Windows user since 1999, and even building a custom Windows PC in 2001, I switched to Linux in early 2009 with OpenSuse (I can't remember the version number, but it was the last release that had a paid box-set available to purchase from Amazon). I switched to Ubuntu later that year with 10.04 LTS, which I used for quite some time before getting a new Windows 7 laptop in late 2010 – which even got upgraded to Windows 8 Pro and then Windows 8.1 Pro albeit with a brief three-month stint of using Ubuntu again.

In January 2014, I got an Acer C720 Chromebook which was my first encounter with Chrome OS, even though I'd been using Chromium OS from December 27th of 2013. It was my first experience of buying a Laptop that comes with a Linux OS as standard – even though Chrome OS isn't a true Linux Distribution. I've been running it alongside my Toshiba laptop ever since I got it, until May 2, when I decided to dual-boot my Toshiba laptop with Ubuntu 14.04 and Windows 8.1 Pro.

The initial install of Ubuntu 14.04 went well – albeit by using an initial install of Ubuntu 10.10 followed by an Upgrade to 14.04 via the Update Manager. I set up a root partition of 232.9 GB for the Ubuntu install, a home partition of 2229.2 GB for my data, a logical partition of 4.0GB for the Linux swap partition, and the remaining 465.8 GB given over to a single partition for Windows 8.1 Pro.

After I got Ubuntu 14.04 updated, I booted from a Windows 7 Ultimate DVD, installed Windows 7, and then installed Windows 8.1 Pro from a 32GB Samsung SD card.

Once I'd got Windows 8.1 Pro installed and set up, I shutdown the Toshiba laptop for the night. Then on May 3, I started up the Toshiba laptop. Instead of being greeted by the usual Grub boot screen with OS selection, I was greeted by the Windows 8 splash screen, and basically a Windows laptop with an unusable set of Ubuntu 14.04 partitions.

After downloading the Ubuntu 15.04 64-bit ISO from the Ubuntu website, and the USB Media creation tool that the Ubuntu site links to, I created a USB installer for Ubuntu 15.04 and Upgrade installed Ubuntu 15.04 onto the Ubuntu root partition. I'm now left with a fully functioning Ubuntu/Windows dual-boot laptop.

After the trouble of getting a Windows/Linux dual-boot system running, in hindsight I might have been better off installing Windows first, and setting up the Windows and Linux partitions from within the Windows installer, and then installing Ubuntu second – to ensure that everything worked correctly when done. However, this is the first time that I've run a dual-boot Windows/Linux system and my only previous dual-boot setup was with Windows 8.1 Pro/Windows 10 Technical Preview.

After learning something new about how to set up a dual-boot Windows/Linux system, I'm now looking forward to using Ubuntu again, and it might even be the kick I need to actually start producing content for my YouTube Channel as there are plenty of professional quality video editors for Linux that might work better on the hardware of the laptop than what Adobe Premier Pro does in Windows. Using the YouTube Capture App for IOS is pretty basic and OK for stitching together several clips. However, using the YouTube Creator Studio from within the YouTube website isn't as powerful as proper editing on a computer and then uploading to YouTube.

Review - KDE Plasma 5

by Alan Ward

First steps with Kubuntu 15.04 and Plasma 5

This April's version of Kubuntu is the first to come by default with the new KDE desktop manager version 5, based on the Plasma 5 user interface toolkit. Previously, in 14.10, two separate desktop environments were offered; the more classical KDE 4 for users more interested in stability, and KDE 5 for early adopters and testers.

KDE has long been known both for its quality graphics, and for its ease of use and configuration. Unfortunately, version 4 was a bit of a disappointment for many. At the beginning, it did tend to consume perhaps more CPU (and GPU) processing than would be reasonable, much RAM, and as a result gave an impression of being generally buggy and working rather slowly. This scared away many, such as myself, who appreciated a well-drawn desktop environment such as KDE version 3 in the past, but also get to need some work done, quickly and on time.

The impression we have is that the KDE project team members have been hard at work making the KDE desktop as a whole much more responsive, and ironing out the various defects that plagued original releases of KDE 4; it certainly seems worth testing out the new KDE, and seeing what progress has been made on this front, and whether the new Kubuntu 15.04 can be used today as a general purpose desktop distribution as it has in the past.

KDE versions 3 and 4

The last version of Kubuntu that came with KDE version 3 (to be precise, 3.5) was Kubuntu 8.04, considered by many as a fairly advanced and stable desktop for the time.

The desktop environment was organized in a way that users of other operating systems would find familiar, with a task bar holding an application menu, direct access links to programs, and a notification area – the whole located at the places one would expect, at the bottom of the screen. Graphical elements' design, such as the icons, was bright and cheerful as well as informative. The icon representing a USB drive was a beautiful piece of work.

Kubuntu 8.04 actually came in two versions, one with KDE 3.5 and the other with the new KDE 4.0 desktop. Then, in late 2008, out came Kubuntu 8.10 with KDE 4.1 as its one and only desktop environment.

It was not much of a success. On the one hand, users were happy to see the KDE project not resting on the laurels earned with KDE 3, but, on the other, the general consensus was that version 4.0 was not yet quite ready for general use. Specific criticisms were levelled at the desktop's perceived slowness, and the fact that some bugs had not yet been ironed out.

Perhaps if KDE 4 had been released as beta software, still maintaining the version 3.5 desktop as an option, critical reception would have been better. In their defense, the KDE project members may have considered that working on a single version of their desktop manager was a better use of limited resources.

KDE 4 introduced a new concept of working with the desktop. Although some elements were distributed in much the same way -- the taskbar, for example, was remarkably similar -- the desktop itself was converted from a mere backdrop that supported icons, to a more global work area in which active elements, desktop widgets, could be placed.

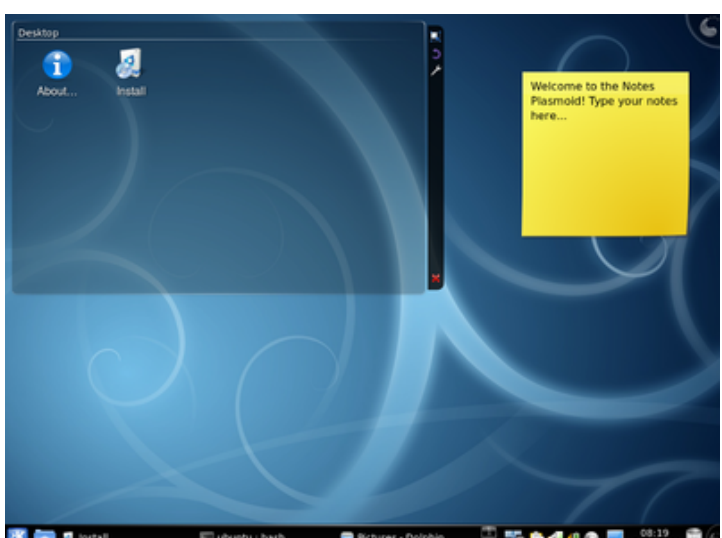

Furthermore, activities were introduced. These worked as separate desktops the user could switch among. Program windows of non-active activities would not interfere with the active pane, but run silently in the background until that activity was activated (brought to the front). Users could organize their workflow across different activities: one for work, another for play, etc.

This is a different concept than virtual desktops -- each activity could contain its own virtual desktops -- although in practice both mechanisms can be used for the same purpose of organizing windows.

This is also when the “Plasma” terminology was introduced, to refer to KDE’s user interface technology as a whole.

The interested reader can still access the original distribution ISO images at http://old-releases.ubuntu.com/releases/kubuntu/. However, it is important to stress that these distributions are no longer maintained. Their use cannot be recommended in production environments, which naturally does not preclude testing on a spare (or virtual) computer if so inclined. When I tested them, Kubuntu 8.10 made itself noteworthy, among other things, by keeping one of the CPU cores (of an Intel Core i5) pegged at near 100% capacity - thus generating much heat and fan noise. This is the type of “feature” that early adopters criticized - and that has been slowly solved over the different versions.

In any case, the KDE 3 series still lives on as the Trinity project (https://www.trinitydesktop.org/), maintained by project members who saw KDE 4 as a step backwards in some areas. This just goes to show the beauty of open-source software: if you do not agree with the direction a specific project is going, just fork it and roll your own. Unfortunately, the Trinity packages are not directly

available within the Kubuntu repositories. For those interested, detailed instructions on converting your existing setup to Trinity can be found here: https://wiki.trinitydesktop.org/UbuntuInstall.

On to Plasma 5

Working with Plasma 5 should not give many surprises to users used to running KDE 4. KDE version 4.14 in Kubuntu 14.10 is much more refined -- and stable! -- than the original 4.0, and the new Plasma 5 environment seems to draw heavily on this now mature version of KDE 4.

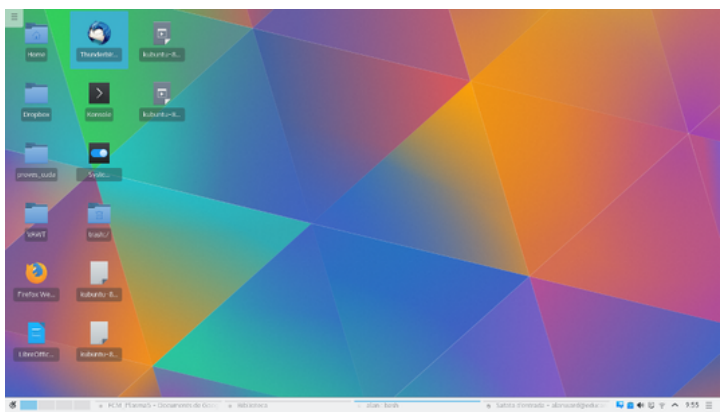

The first impression we have when running Kubuntu 15.04 Vivid Vervet is that of color. The default desktop has left the traditional KDE world of blues, and gone towards a more harlequin-inspired set of colors.

Application and other icons have evolved towards a flatter presentation, leaving behind the previous skeuomorphic tendency of icons with a 3D effect in their graphism. This is exactly the same tendency that is seen in Google's web design, Microsoft's Windows 8 Metro environment, and Apple's IOS 7 and OS-X 10 Yosemite. In the beta versions of 15.04, some of the icons had not yet been updated, as is the case with Thunderbird in the screenshot.

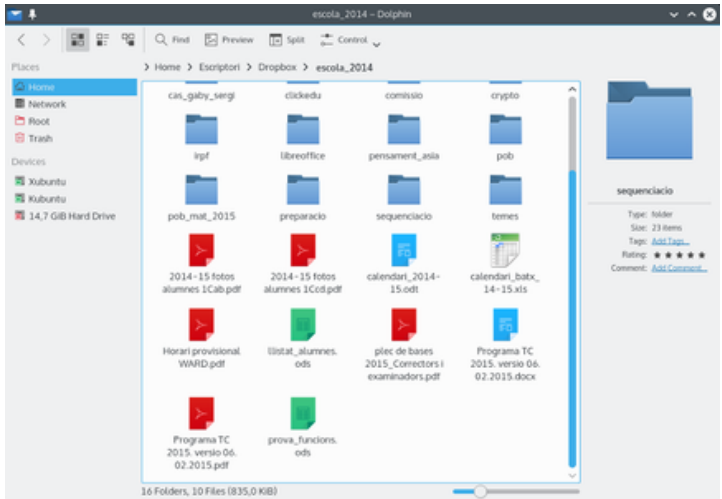

The new design is pervasive across the desktop environment. Besides the main (large) icon set used for files, the status bar, menu bars, the Dolphin file manager's lateral bar, and various dialog windows, all share a second, understated and smaller version of the new icon set. This mini-icon set is based on pure line art, in black for informative icons and red for alerts, on a gray background. The lack of details combined with a very small base icon style can make for difficult reading at times. Perhaps the designers were thinking of higher resolution viewing screens, where the thin lines may be better rendered.

A window theme called Breeze replaces KDE 4's Oxygen and replaces the light or dark gray gradients we know with wide, uniform dark window bars. Widget buttons show new OK and Cancel buttons with revised graphics based on thin lines, also similar in a way to Apple's newest offerings.

Notifications are rather low-key, appearing briefly in the low right-hand corner. Each individual notification can be managed on a granular basis: we can choose for each action whether we prefer a message, a sound, both or

neither.

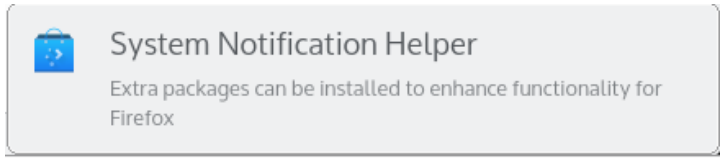

Clicking on the notification message gets us, in many cases, to the relevant program to solve the problem; for example, from an alert concerning lacking packages, we can run the relevant installer directly. After installing the public release version of Kubuntu 15.04, the updater came up immediately with a request to install a security update. This is not quite rocket science as far as user interface design is concerned, but is certainly helpful in day-to-day system management.

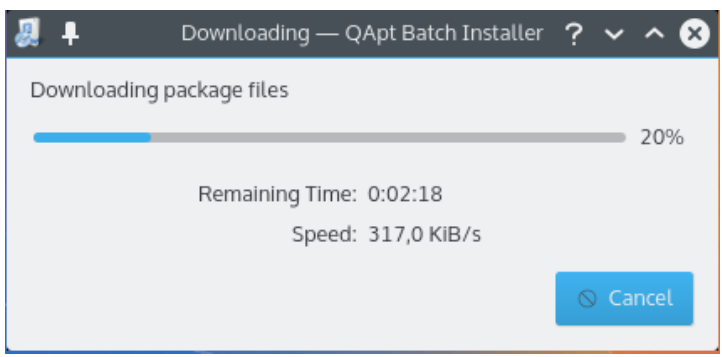

The driver manager is also there, ready to inform the user about the availability of specific drivers for the hardware.

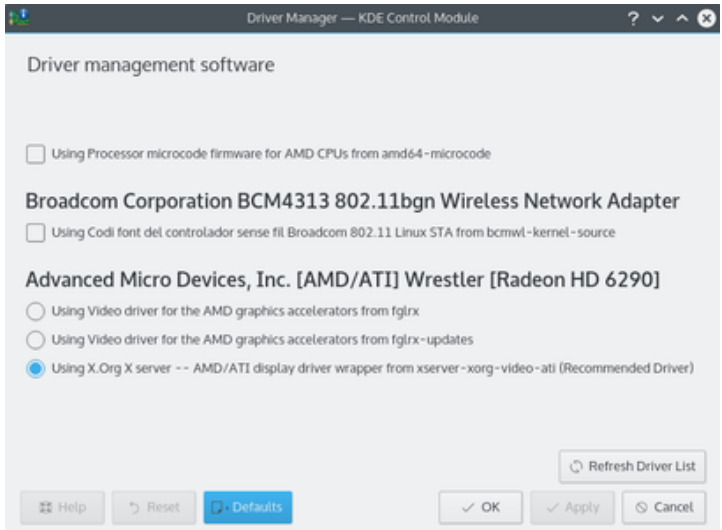
Driver Manager — KDE Control Module
Driver management software
Using Processor microcode firmware for AMD CPUs from amd64-microcode
Broadcom Corporation BCM4313 802.11bgn Wireless Network Adapter
Using Codi font del controlador sense fil Broadcom 802.11 Linux STA from bcmwl-kernel-source
Advanced Micro Devices, Inc. [AMD/ATI] Wrestler [Radeon HD 6290]
Using Video driver for the AMD graphics accelerators from fglrx
Using Video driver for the AMD graphics accelerators from fglrx-updates
Using X.Org X server -- AMD/ATI display driver wrapper from xserver-xorg-video-ati (Recommended Driver)
Refresh Driver List
Help
Reset
Defaults
OK
Apply
Cancel

Activities and virtual desktops work in the same way as KDE 4. So do the various types of desktop organization ("workspace type" in KDE terminology). KDE 5 includes two types of workspace that can be configured from the desktop settings window (right click on the desktop):

- A "standard" desktop view, in which no elements are actually located on the screen.
- The "folder" view, in which the contents of directory /home/<username>/Desktop are reflected on the desktop. This is the view that best corresponds to other operating systems.

Previously, KDE 4 also had two further view types:

- A "newspaper" layout, that organized windows and other desktop widgets in columns.
- The "netbook" view, with a main menu that covered the entire desktop, much like tablet and phone user interfaces.

These seem to have disappeared in KDE 5 for some reason; they can no longer be downloaded from the

repositories, and the option in System settings > Workspace Behavior > Workspace > Workspace Type has disappeared. We are still not sure whether this is just a temporary setback that will be corrected later on, or a conscious decision to simplify the number of options in this environment.

Perhaps the most noteworthy feature of Plasma 5 is the lack of specific new additions. Some changes have been made, such as the login display manager (the new SDDM replaced KDM), or the Baloo file indexer replacing Nepomuk. As for the rest, the systems settings configuration panel, the file manager, and other components of the desktop environment work much in the same way as KDE 4. In this sense, KDE 5 can be seen more as a continuation of KDE 4 than a revolution. There are fewer differences between them from a user's perspective than between versions 3 and 4.

Speed is important

As stated above, the final releases of KDE version 4 already gave better performance figures than previously. This tendency is confirmed by Plasma 5. We have tested the new desktop on two platforms. One was an Intel Core i5 laptop with an SSD hard drive, while the other was one of the original Acer Aspire netbooks, equipped with an Intel Atom 1.6 GHz processor, 1 GByte of RAM, and a spinning drive.

Performance of Kubuntu 15.04 was good on the Core i5, as expected. The system comes up quite sprightly, and after some rather intense web navigation and the use of programs, reports a mere 1.2 GBytes RAM occupation (out of a total of 4 GBytes). The dedicated nVidia graphics processor on this laptop also helps get the most out of KDE's desktop effects.

On the other hand, I approached the small Acer with a bit of trepidation. I had abandoned all ideas of using a complete desktop manager with all the graphical bells and whistles on this particular piece of equipment; its purview was restricted to lightweights such as Xubuntu's XFCE and MATE - where it did quite well, indeed.

Rather to my surprise, the Kubuntu 15.04 (i386) software + small Acer hardware combination booted up in 1 minute 46 seconds. This may seem a tad slow at first sight, but is actually quite correct given the age of this (rather low-range) piece of hardware. Obviously, boot-up times are at a much faster 40 seconds on the Core i5 due to the combination of a more powerful processor and, most of all, the SSD hard drive.

Sound and graphics work well out of the box on the Acer, as does WiFi. I was able to navigate the web and even view videos on Youtube with no hassle. Video reproduction is a bit choppy for any resolution over 400 pixels, which is perhaps excusable for the available

hardware. The hardware manager immediately came up, and proposed the installation of the single proprietary driver applicable for this particular laptop (Intel CPU microcode).

A specific tweak applied to this computer, that can be recommended for computers with slower hard drives, is to disable file indexing. The Baloo indexer can be configured to index only some volumes or deactivated altogether in the System settings > Desktop effects control panel. Otherwise, the user can expect some heavy disk usage, at least the first times the system is booted up. On the Core i5, Baloo's activity resulted in one of the processor cores getting pegged at 100% capacity for some time, as it digested the contents of the hard drive.

The overall impression is that much attention has been given to staying more-or-less within the general desktop design parameters initially drawn up for KDE 4, while increasing actual usability on a wider range of hardware platforms.

Not all is rosy

Unfortunately, there is still some scope left for making improvements. There was a noticeable leap in quality from Kubuntu 15.04 Beta 1 to Beta 2, and more bugs have been ironed out in the final release to the public. The apport bug notification application had nothing to do after setting up and booting the final release, unlike both betas. Perhaps it may be useful to remember that beta versions are not finished software releases, but must be considered as work-in-progress at best, and downright buggy at worst. However, by making them available to early adopters, more people will use the new software in a variety of situations and experience possible bugs. By reporting them to the Kubuntu or KDE projects, as appropriate, developers can get on top of them before the

final release is offered up to the general public.

The beta versions of Kubuntu 15.04 did contain quite a lot of bugs. For example, the installer simply locked for some testers, while, in my case, I got a system that went into sleep recurringly every two minutes while the installer was doing its job. Once the system was up and running, upgrading the kernel to version 3.19 did not work well for some, needing a reboot back into the previous version, 3.16.

A noticeable quirk is that the kwallet application reported needing to migrate from a previous file layout to a more modern version. This happened to me even when performing a fresh install with the public release, which was a bit weird.

Perhaps the most obvious defect for the general user concerns the notification area in the menu bar. Although general operation is similar to KDE 4's offering, the internal mechanisms seem to have changed a bit under the hood. This results in not playing well with many external software applications. Dropbox works well in the background, but with no status icon to inform the user about its state. It is difficult to know precisely what the program is doing, which can be a bit of a hassle - especially at the beginning, just after installation, when Dropbox needs to download some binary blobs and index files on the hard drive. Some versions of Plex also seem to have problems, refusing to start and complaining about not being able to access the notification area.

Some computers, such as a low-cost Acer AO-722 notebook, had issues with the screen flickering and window borders disappearing as soon as multiple virtual desktops were enabled. This seems to depend on the graphical hardware present, and can be cured by installing specific graphical drivers or, obviously, by

deactivating virtual desktops.

More serious was a problem in Beta 2 that left the user looking at a blank screen. Rebooting did not seem to help, either. This seems to be related to a problem with the configuration files, that moved from directory ~/.kde to ~/.config/kde and ~/.config/plasma. The solution given in https://www.kubuntuforums.net/showthread.php?67234-Black-screen-with-Plasma is simply to remove these directories (be careful!), and reboot the computer. This wipes out configuration items such as WiFi networks but makes the desktop come up once more.

Once again, let us remember that detecting such problems is precisely why beta versions exist. But there do seem to have been a lot of them in this version of Kubuntu, most of which were corrected in the public release of April 23.

To end on a more positive note, some of us were wondering if the -admittedly major- system component change towards systemd would affect system usability. I am happy to report this is not the case for Kubuntu 15.04, at least as far as general desktop or laptop use is concerned. I would be more concerned from the standpoint of server administration, but, on the other hand, a KDE desktop would probably not be the system of choice by the people concerned, for that specific type of workload.

Aside from the minor difficulties experienced, I would like to end with a general comment on this particular version of *buntu 15.04: that the user experience is just about what some of us have been hoping for from the Kubuntu project for some time now. We have a nice, colorful desktop with a lot of helpful automated actions going on under the hood to take the hassle out of basic systems administration. It is safe to say that this version should gain acceptance from users, and perhaps not only among

power-users and staunch supporters of KDE but also among the more volatile regular desktop switchers such as yours truly. This distribution also impresses favorably for its speed on low-end and elderly computers, though it should be stressed that in such a case reducing the amount of desktop effects active can help avoid processor overload. In any case, as usual with KDE, if a processor with decent speed and a fast -- preferably SSD -- hard drive are available, the desktop environment can make good use of them.

Author biography: Alan teaches computer science at Escola Andorrana de Batxillerat (high-school). He has previously given GNU/Linux courses at the University of Andorra and taught GNU/Linux systems administration at the Open University of Catalunya (UOC).

Q&A 97

By Gord Campbell

Q How can I get my Nvidia GTX 960 to work properly with Ubuntu?

A Google this: gtx 960 ubuntuforums

The first result should point to a post which says "SOLVED". The procedure can probably help people with other recent Nvidia graphics cards.

Q I'm remotely mounting a shared Windows 8 partition from Ubuntu 14.04.2 using an entry in fstab. When I look at the file permissions, it looks as if I should be able to write to the partition. However, when I try to write to the folder I get "Permission denied."

A (Thanks to Mark Phelps in the Ubuntu Forums) Check in Win8 to ensure that you have disabled FastStartup -- that is NOT the same as Fast Boot. Also, when you leave Win8, make sure you choose ShutDown, not Restart, because the latter automatically enables hibernation.

Q I am setting up a new web/game server and want to have remote access, but we want it more secure than SSH. What should we do?

A (Thanks to TheFu in the Ubuntu Forums) I don't know of anything more secure than SSH using ssh-keys. Never allow passwords for Internet connected systems; that is TOTAL FAIL.

Lots of ways to lock down SSH further:
- install fail2ban (blocks brute force attacks)
- block all external IPs except from the people who should have access
- use an iptables rule that slows down access from IPs with multiple login failures.

If you want more security with ssh, you could force two-factor authentication with something like Google-authenticator or Ubikey.

Q How can I have a plain background in Xubuntu 15.04?

A In Settings, Desktop, there is a drop-down called "Style." Select the top item, "None." The line below that lets you pick the color for your plain desktop.

Q I have an ATI Radeon HD 2400XT, and, when I installed Ubuntu 14.04, the graphics card driver was installed automatically. Did Ubuntu install the latest ATI graphics driver or is it a generic driver?

A (Thanks to QIII in the Ubuntu Forums) That is an older card and it no longer enjoys Linux support from AMD. The legacy driver will not work with the current version of X Server. When you installed Ubuntu, the default open source Radeon driver was installed -- otherwise you would not be seeing anything. That is the driver you will need to use.

Q I need to mount my laptop upside down. With the Nvidia 331.113 drivers my games perform well but I cannot rotate my screen.

A Run Nvidia X Server Settings. Select "X Server Display Configuration" on the left. The fourth item on the right is "orientation".

Q I added LibreOffice to a fresh install of Xubuntu 15.04, but the spell check did not work.

A Run Language Support and follow its suggestion to complete the installation. Restart LibreOffice.

Q I am interested in donating some old computers.

A (Thanks to Veddox in the Ubuntu Forums) That is a good idea, if they are still reasonably good. Speaking as a teacher at a school in Africa, we sometimes get old computers donated by somebody in the West. A few are still usable, but all too often they have about 128MB RAM, broken touchpads and an Intel Pentium from 1999. I'm sure the people who send these out do so with the best intent, but they don't help, we can't do anything with them. In short: don't give away junk. (If they have enough power to run Windows 7 reasonably smoothly, or Ubuntu or Linux Mint, they are very much welcome. But if you can't even start AbiWord on them without crashing them, no thank you!)

Top questions at Askubuntu

- How is 'rm -rf /' able to delete all files in the system? http://goo.gl/z3bbEn
- What does this 'apt-get remove' instruction do?

http://goo.gl/UO621J

- How can I create a video file from a set of jpg images?
  http://goo.gl/Fdspz8

- How do I prevent one user in particular from accessing my home directory?
  http://goo.gl/M7qOpi

- "mv file" and now it's missing
  http://goo.gl/qSBKH2

- Can I emulate graphical DOS applications inside a TTY?
  http://goo.gl/rtO7HF

- How does Ubuntu know the make and model of my laptop?
  http://goo.gl/ZrQcWS

- How to install Visual Studio Code on Ubuntu?
  http://goo.gl/pl0Gxw

- How can I merge files on a line by line basis?
  http://goo.gl/3sRV81

Tips and Techniques

A new member of the family

Ubuntu Mate is now an official community distribution in the Ubuntu family. I have an Acer Aspire One netbook which wasn't up to the task of running Ubuntu Kylin, and Mate has a reputation of being fairly lightweight, so I decided to have a go.

What do I want from a computer? Remote desktop host

(x11vnc) and client (KRDC), file server (samba) and print server, web browser (Chrome), media player (VLC), productivity applications (LibreOffice), webcam recording (guvcview), and system monitor (conky). Everything worked just fine. It felt a bit slower than Xubuntu on the same computer, but not by much.

When you start Mate, it has panels at the top and bottom of the screen. With the tiny screen of the netbook, only 600 pixels high, one panel had to go. A few minutes of work did the job, although there is still room for improvement. (I lost network manager and the volume control in the shuffle, but they are both available in Settings.)

After installing Synaptic Package Manager, I searched for "samba" and installed everything which looked relevant. Five minutes later, I had a shared folder which allowed everybody to read and write files. (This was the area where Mate had an advantage over Xubuntu.) The netbook isn't very powerful, but it has a 160 GB hard drive, which is enough to be useful for backing up current projects. (I may even set up Own Cloud, as per the last issue of FCM.)

I plugged in my Brother laser printer, and eventually got it shared. The trick: in Printers, Server Settings, enable "Publish shared printers connected to this system."

Six years ago, this netbook was the slowest computer on the market, so I'm not planning to make heavy use of it. Even so, with Mate it is a usable ultra-portable, small file server or even a remote camera.

On the other hand, the advantages over Xubuntu were very limited.

Security

by Kevin O'Brien

SSH options, Port Forwarding over SSH, Keepalives

The ssh command has a number of options, and I don't plan to cover all of them. Even the SSH documentation warns against the use of some of them, suggesting they are only for real experts. But I want to mention the ones that I think you will find important. These options take the form of switches in the command:

- -1 : Forces the connection to use SSH v.1 protocol only. The question here is why would you want to do that if you have SSH v.2 available. It is a real improvement, after all.
- -2 : Forces the connection to use SSH v.2 protocol only.
- -4 : Forces ssh to use IPv4 addresses only.
- -6 : Forces ssh to use IPv6 addresses only.
- -b : Bind address. Useful for machines that have two IP addresses, such as systems with two NICs. This tells SSH which IP address on the local machine to use for the connection.
- -L : Specifies that the given port on the local (client) host is to be forwarded to the given host and port on the remote side. This works by allocating a socket to listen to port on the local side, optionally bound to the specified bind_address.
- -p : Port to connect to on the remote host. This can be specified on a per-host basis in the configuration file.
- -R : Specifies that the given port on the remote (server) host is to be forwarded to the given host and port on the local side. This works by allocating

a socket to listen to port on the remote side, and whenever a connection is made to this port, the connection is forwarded over the secure channel, and a connection is made to host port hostport from the local machine.
- -v : Verbose mode. This shows all commands and replies, and is useful for debugging.
- -W : Requests that standard input and output on the client be forwarded to host on port over the secure channel. Works with v.2 only.
- -X : Enables X11 forwarding. But note that this can open a vulnerability.

Port Forwarding

One of the handy things you can do, and something useful for tunneling. is port forwarding over SSH. The basic idea is to connect via ssh to a remote machine, and ask it to send something to a specific port other than the default port. The basic way you do this is to use the SSH command with the appropriate flags, -L and -R, which, not surprisingly, stand for Local and Remote. You need to specify the port you want to use, and what will be forwarded to it.

- Local Port Forwarding – This takes a port on your local machine and forwards it to a specified port on the server. So you can make a request on a local port like 7280 on address 127.0.0.1, and your SSH client would intercept that call and send it to port 119 on the server. Then you would have a secure connection to get whatever port 119 is configured to serve (typically Usenet traffic, but this is just an example). So you use this to configure your newsgroup client to securely grab messages from a public server, assuming it allows SSH connections.
- Remote Port Forwarding – This is the reverse of Local Port Forwarding. Here, the idea is to specify a

port on the remote server and have it forwarded to your local server. This is not very common, and you may never need to do this. Essentially, all traffic coming in to the server on the specified port would then be forwarded to your local machine.
- Dynamic Port Forwarding – This creates a SOCKS proxy and is not restricted to one port or one type of traffic.

Local Port Forwarding

Suppose you are at work (or school), and you just cannot bear to miss out on your Facebook stream. But there's a filter stopping you from accessing the site. However, for the sake of argument, you could create an SSH connection to a server outside the network (which could be your computer at home). You could then do something clever using Local Port Forwarding. Create a connection as follows:

```
ssh -L 7280:facebook.com:80 address of home machine
```

Now, your home machine does need to have a public IP address, or you would need to set up your router to forward the traffic, for this to get through. Once you have done this, you would open your browser and set it to go to http://localhost:7280, and traffic would then flow to your home machine, and from there to Facebook. You can now browse to your heart's content on your Facebook stream. Of course, this also illustrates why network admins might want to shut down SSH traffic, which they might do by blocking any outbound traffic going to port 22 (the default SSH port). And you could then try changing the default port on your home SSH server to something other than port 22, and then the admins could do deep packet inspection, and so on.

But SSH Port Forwarding is not just a matter of a security

breach in the making, it can be used very legitimately in a number of situations. For example, you have a company with a number of geographically dispersed locations. In that case, SSH Port Forwarding would be a very useful way to connect sites to exchange data. You might have a database server that employees might need to connect to, and don't want that traffic flowing through the Internet unsecured. Or perhaps you have set up a server for yourself, such as OwnCloud, and it is in a remote hosting center. Creating an SSH connection and using Port Forwarding might make your data a lot more secure.

Limitations

There are a few things you need to watch out for. One is that not all ports may be available to you. If you are in a Unix-like environment, for instance, port 1024 and all ports below that can only be used by root. But any port above 1024 should be usable by a user with normal privileges as long as no one else is already using it.

The other thing you need to remember is that if the connection is dropped the port forwarding is gone. And, in general, TCP connections are configured to close after a period of inactivity, and on some firewalls that can be as little as 300 seconds (5 minutes). This can be controlled by a rule (or perhaps more than one) in your iptables, or directly by /proc/sys/net/ipv4/tcp*keepalive*time. But, if you want a persistent connection, you need to use a Keep Alive.

Keep Alives

There are two, basically. One is the TCP Keep Alive, which is simple but spoofable, and the other is the SSH keepalive, also called serveralive. Serveralive messages travel through the encrypted connection between you and the server, and thus cannot be spoofed. Assuming security

is your reason for creating SSH connections, then it would be more secure to use serveralive messages, though I expect using TCP keepalives is far from the worst thing that could happen. In Linux, you can set this up either for everyone (if you have root privileges), or just for yourself, by editing the appropriate config file.

For everyone – edit /etc/ssh/ssh_config and insert the following:

```
Host *
ServerAliveInterval 300
ServerAliveCountMax 2
```

For just you, edit ~/.ssh/config, and add the above code to it.

The first line, ServerAliveInterval, specifies how often a null packet should be sent to the server to keep the connection alive. However, sometimes the server may go off or drop the connection, so the second line specifies how many times you should send a packet without getting a response. The setting I have shown will send a packet, and if no response is received, it will send a second packet 300 seconds later. If no response is received to the second consecutive packet the connection will be dropped by your client.

On Windows, assuming you are using PuTTY, there is a good explanation at http://blog.hazaveh.net/2013/10/keep-ssh-session-alive-in-putty/, but essentially you go to Connection, and then on the right under Sending of null packets to keep session active, you can set Seconds between keepalives (0 to turn off) to 300 seconds to get a similar result.

Once you understand Port Forwarding and Keepalives, you are most of the way to SSH tunneling,

Ubuntu Phones

Meizu launches the Ubuntu MX4

In the first of a series of launches, Meizu has announced that the Ubuntu MX4, which will be sold across Europe soon, is immediately available to developers in China. The move kicks off Meizu's push towards strengthening the Ubuntu ecosystem in China, ahead of launching the Chinese Edition to its user base later this year.

Meizu is one of the most popular high-end smartphone brands in China, recognized for its innovative and elegant product design and for pioneering the user centric distribution model that has become the standard for several Chinese manufacturers.

The Ubuntu MX4 launch also coincides with the final months of the China Mobile Ubuntu Developer Contest, which is already seeing thousands of developers participating in onsite and online training as well as hackathons at various universities across China.

Ubuntu 14.10 (r22)

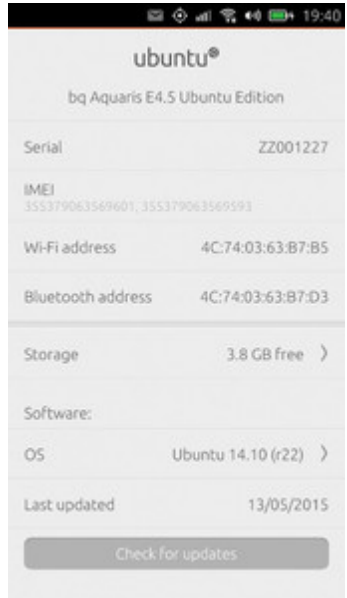

Brief list of changes in stable's OTA-3.5

- Fixed some issues with mobile carrier data not connecting when leaving Wifi
- Fixed calendar sync
- Fixed wrong hang-up of calls when there is a call waiting
- Fixed ubuntu-keyboard crash
- Disable upstart logging to fix cases of reboot loops
- Fixed listing click updates for frameworks newer than supported on the device
- Improve location accuracy
- Fixed routing problems after switching from Wifi to mobile carrier data
- Fixed multiple telepathy accounts from being created
- Fixed store scope updates causing the scope disappearing from favorites

- Fixed airplane mode causing battery drain
- Fixed issues with enabling location
- Fixed issues with non-releasing wakelocks causing phone not to suspend
- Fixed indicator-network crash
- Fixed location settings persistance after reboot

Tuxidermy
AND I WAS TELLING OUR LITTLE BARTENDER HERE I'M A FREAKING GENIUS, BABE!
JEEZ.
I CAN, LIKE, BUY AND SELL THE WHOLE WORLD IF I WANT TO, YOU SEE? I'M LIKE A KING, A GOD...
A JERK...
SO, HOW ABOUT WE START MAKING OUT ALREADY, HUH? COME ON, GIVE YOUR BOSS A KISS.
DON'T YOU DARE!
SUDO RM -R STUPIDMORON
WAIT. WHAT'S THAT?
CLICK!
THERE YOU GO, MISS.
WOW! THANK YOU! WHAT'S DOWN THERE, HANDSOME? IT LOOKS LIKE A LONG FALL.
ABANDONWARE.
NICE. POUR ME ANOTHER ONE, WILL YOU?
IT'S SO DARK IN HERE!

Ubuntu Games
This War of Mine

by Oscar Rivera

Most of the time when we play a war game, we are either a foot soldier in a first-person shooter, we drive a combat vehicle (like in War Thunder reviewed in FCM#93), or we are the prime decision-maker in war strategy games like Empire: Total War. The video-game This War of Mine fits none of those scenarios. In fact, in This War of Mine, you are not really even fighting a specific enemy per-se, except perhaps for the war itself. Erase every idea you may have about a war game, and instead imagine yourself in any war that's ever been – as an ordinary civilian trying to stay alive. This War of Mine is a war, survival & strategy game where you're just trying to stay alive as a civilian in a war that you didn't ask for and you wish would just end. “In war, not everybody is a soldier” is the quote used to promote the indie game This War of Mine. That phrase about sums it up as to the content of the game. This War of Mine was released for Linux, Mac & Windows on November 2014 by Poland-based 11-Bit Studios.

Game Synopsis

This War of Mine depicts the gruesome reality of the horrors that the average person must endure when a war happens to knock at your front door. It's most certainly a war-survival game, but it is also definitely a strategy game where you must plan out your course of action from the get-go. I will add right now that This War of Mine is NOT for the faint of heart. It's not your typical combat game where you'll see bloody, gory graphics, but instead the warning I've issued is due to the moral choices you'll be

forced to make in order to survive and the unimaginable events you'll witness before the war ends. I've never had to endure a war first-hand but playing this game reminded me of when I was reading Viktor Frankl's account of surviving in a WWII concentration camp in his book Man's Search for Meaning. You will definitely need to make some tough choices in this game – which is one of the reasons it's gotten such rave reviews since its release. The game itself is very well polished and, out of all indie games I've played in recent times, it's without a doubt the most accomplished and best all round. I have not encountered any bugs yet, and, from what I've read in various forums, I don't expect to find any bugs at all.

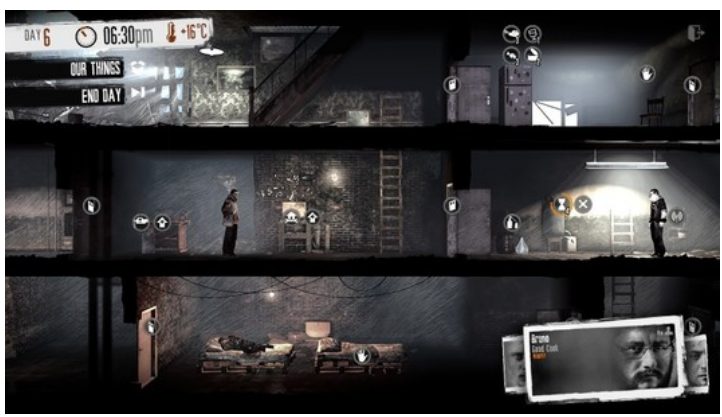

Playing the game

You start out the game with a team of survivors, usually three or four, and, as the days pass, there's a chance that at least one other person will knock at your door and ask to join you; it's completely up to you whether you accept the other person or leave them outside to fend for themselves. Not having a tutorial at all, the first couple of times you play, it may be a sort of trial-and-error learning experience, but, after a couple of tries, you should have a handle on what you're doing and begin to plan a well adjusted strategy. My first time playing I actually made it

to nine days before my first character was killed.

The only, and I do mean the ONLY, downside to this game is that you don't have the luxury of saving your progress like you do in other games. The only options available to you when you play are “Continue” or “Another Try” which should be self-explanatory.

Your survivors are found in a decrepit building that has plenty of items scattered around to get you started. The first order of business is to send your survivors to scavenge whatever they can find in your building. Unfortunately, there will be places that will be unreachable at first until you obtain either a lock-pick or a crowbar, both of which can be built, or traded from someone else. Having gathered as much as you can, you then face the choice of what to build with your materials. In order to have my players well rested, I like to build two beds for them, though at first one will suffice. You may also want to build a metal workshop so you can build your crowbar and other metal needs. Ultimately, all you really need to survive is rest (on a bed is preferable), food, water, and meds. If you have a character with addictions such as coffee or cigarettes, you will also need to provide those, otherwise your character will become depressed and end up broken, catatonic, and/or leave your team.

Building a heater is essential when winter comes, an animal catcher will help you catch rats to eat, a distillery will provide you with alcohol for trading, and the list goes on. In the daytime, while staying in your building, it's your job and responsibility to upgrade your living arrangements, cook food, gather water, make alcohol, upgrade your tools, grow herbs, make cigarettes, meds, etc. At night, you send one of your survivors out to the fictional city of Pogoren so they can scavenge for more items. The next day is all about further upgrading your living quarters with the items you have found at night in other buildings around the city. The real danger shows its ugly face at night when you're scavenging around the city. Some places are relatively safe to visit; however, others are not. There are places you can go to where armed assailants are willing to kill you at first sight. Other places have inhabitants still occupying their original abode, and may or may not be willing to trade with you. Regardless of whether the trade happens, if you happen to take any items (especially food or meds) that belong to them, your survivors will become sad because they know that those people's lives might have been shortened thanks to you robbing them. However, taking items from criminals is OK and has no negative feelings associated with it.

For defense there are a number of weapons that can be used throughout the game, such as knives, crowbars, shotguns, etc. What I've found best is not to kill unless you absolutely have to, and, if you do, it's best to kill with a knife or another silent weapon instead of something loud like a shotgun that will alert other criminals of your whereabouts.

Playable Characters

There are different permutations of people with whom you can start out the game. There are a total of 10 playable characters, and each has strengths and weaknesses. Some combinations of people are better than others depending on what strategic approach you will take. For example, Bruno is a good cook, and your whole team benefits from him in that he uses less resources when cooking food, collecting rainwater, distilling alcohol, or producing meds or cigarettes. Katia is good at bargaining, and whenever you have to barter with someone, she'll get you the best deal. However, for scavenging, the best is no doubt Marko as he can carry up to 15 items in his backpack, and he’s not slowed down when the load is heavy.

The War is Over

You eventually beat the game by surviving anywhere between 20-50 days. I've heard of some people beating after only 22 days, and yet there are others who had to reach 47 before beating it. There are different scenarios played out as your game progresses which are decided depending on your in-game choices as well as just pure random luck. For example, your game could begin in the middle of winter, which would make you start the game already owning at least one heater. On the flip-side, your game could begin before winter approaches which means you have no heater but when winter arrives you must build a heater, otherwise you'll freeze to death... literally. I've had a couple of characters actually freeze to death. I've also had characters die of hunger, illness, or even untreated injuries. The quick way to die is getting killed while scavenging throughout the city. Your characters are also prone to getting depressed and contemplating suicide if times are really tough. In short, it is as if you really are trying to survive a war that just landed in your backyard. Like I've previously stated, it is the decisions you make that ultimately determine how likely it is that your characters will survive in this war-torn state.

Installing This War of Mine

You can install This War of Mine in one of several ways. If you want all of your proceeds to go to 11Bit-Studios, then you can get the game from the 11bitstudios website. Otherwise, you can get it from either Steam or www.humblebundle.com. The download and installation is pretty straightforward. I got mine from Steam, and, after purchasing it, I simply waited a few minutes for Steam to download and install it in the background while I was busy doing other things. The game currently sells for $19.99, but often it is on sale for lower than that. I got lucky and paid half off the full price when I got it.

Minimum Specs

These are the minimum specifications as taken from the official game website:

OS: Ubuntu 12.04
Processor: Intel Core 2 Duo 2.4 GHz, AMD Athlon X2 2.8 GHz
Memory: 2 GB RAM
Graphics card: Geforce 9600 GS, Radeon HD4000, Shader Model 3.0, 512 MB RAM DirectX: Version 9.0c required

My Custom-made Gaming Rig

AMD FX-6100 3.3GHz CPU (overclocked to 3.5GHz)
Asus M5A97-EVO motherboard
Gigabyte Windforce GeForce GTX 960 graphics card with 346.59 proprietary driver
8GB of Kingston Hyper X RAM & 1TB Seagate Barracuda hard drive
Ubuntu 14.04.2 LTS with Unity desktop

Conclusion

In short, this is by far one of the most entertaining & innovative games I've had the pleasure of playing in a long time. I really enjoyed playing This War of Mine. The game-play is unique and keeps you engaged. The way it stirs up emotions within you is something I've not experienced from a video game before. The charcoal-stylized graphics intertwined with authentic street art pieces made by artists from around the world (included only if you buy the War Child Charity DLC) keeps the game uncomfortably gloomy. The eerie and depressing soundtrack unmistakably enhances the uneasy feeling that permeates throughout the game. I strongly recommend this game.

5 out of 5 stars

Oscar graduated with a music degree from CSUN, is a Music Director/Teacher, software/hardware beta tester, Wikipedia editor, and active member of the Ubuntu community. You can follow him at: www.gplus.to/7bluehand https://twitter.com/7bluehand or email him at www.7bluehand@gmail.com

FLAT
20%
Discount
[PACKT]
PUBLISHING
Learning AndEngine
Martin Varga
Offer Valid till - 21st November 2014
Use Discount Code - gUdb4pzLI
http://bit.ly/1rG1qDE

Contribute to Full Circle Magazine
FULL CIRCLE NEEDS YOU!

A magazine isn't a magazine without articles – and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow *buntu users. Send your articles to: articles@fullcirclemagazine.org.

We are always looking out for new articles to include in Full Circle. For help and advice, please see the Official Full Circle Style Guide: http://url.fullcirclemagazine.org/75d471

Send your comments or Linux experiences to: letters@fullcirclemagazine.org
Hardware/software reviews should be sent to: reviews@fullcirclemagazine.org
Questions for Q&A should go to: questions@fullcirclemagazine.org
Desktop screens should be emailed to: misc@fullcirclemagazine.org
... or you can visit our forum via: www.fullcirclemagazine.org

Full Circle Team:
Editor - Ronnie Tucker - ronnie@fullcirclemagazine.org
Webmaster - Rob Kerfia - admin@fullcirclemagazine.org
Podcast - Les Pounder & Co. - podcast@fullcirclemagazine.org

Editing & Proofreading:
Mike Kennedy
Lucas Westermann
Gord Campbell

Robert Orsino
Josh Hertel
Bert Jerred

Our thanks go out to Canonical, to the many translation teams around the world, and to Thorsten Wilms for the current Full Circle logo.

Getting Full Circle Magazine:

EPUB Format - Most editions of Full Circle have a link to the epub file on the downloads page. If you have any problems with the epub file, you can drop an email to: mobile@fullcirclemagazine.org

Google Currents - Install the Google Currents app on your Android/Apple devices, search for 'full circle' (within the app), and you'll be able to add issues 55+. Or, you can click the links on the FCM download pages.

Ubuntu Software Centre - You can get FCM via the Ubuntu Software Centre: https://apps.ubuntu.com/cat/. Search for 'full circle', choose an issue, and click the download button.

Issuu - You can read Full Circle online via Issuu: http://issuu.com/fullcirclemagazine. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu Linux.